Hierarchical Time-Dependent Oracles*

Spyros Kontogiannis[1,4], Dorothea Wagner[2], and Christos Zaroliagis[3,4]

1 **Comp. Sci. & Eng. Dept., University of Ioannina, Greece** [kontog@cse.uoi.gr]
2 **Karlsruhe Institute of Technology, Germany** [dorothea.wagner@kit.edu]
3 **Comp. Eng. & Inf. Dept., University of Patras, Greece** [zaro@ceid.upatras.gr]
4 **Computer Technology Institute and Press "Diophantus", Greece**

arXiv:1502.05222v4 [cs.DS] 1 Jul 2016

Abstract

We study networks obeying *time-dependent* min-cost path metrics, and present novel oracles for them which *provably* achieve two unique features: (i) *subquadratic* preprocessing time and space, *independent* of the metric's amount of disconcavity; (ii) *sublinear* query time, in either the network size or the actual Dijkstra-Rank of the query at hand.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems; I.1.2 Algorithms 05C85 Graph algorithms; 05C12 Distance in graphs;

Keywords and phrases Time-dependent shortest paths; FIFO property; Distance oracles.

Digital Object Identifier 10.4230/LIPIcs...

1 Introduction

Concurrent technological infrastructures (e.g., road networks, social networks, e-commerce platforms, energy-management systems) are typically of very large scale and impose as a routine task the computation of min-cost paths in real-time, while their characteristics usually evolve with time. The large-scale and real-time response challenges have been addressed in the last 15 years by means of a new algorithmic trend: the provision of *oracles*. That is, data structures created by appropriately selecting precomputed information (summaries) and which subsequently support query algorithms with real-time responses. The quality of an oracle is assessed by its preprocessing space and time requirements, the time-complexity of the query algorithm and the approximation guarantee (stretch). Numerous oracles have been proposed and analyzed (see e.g., [1, 20, 21, 23, 24, 26, 27, 28] and references therein) for large-scale, mostly undirected networks, accompanied by a *static* arc-cost metric. In tandem with oracles, an equally important effort (with similar characteristics) has also emerged in the last 15 years under the tag of *speedup techniques*, for approaches tailored to work extremely well in real-life instances (see e.g., [3] and references therein).

The temporality of the network characteristics is often depicted by some kind of predetermined dependence of the metric on the actual time that each resource is used (e.g., traversal speed in road networks, packet-loss rate in IT networks, arc availability in social networks, etc). Perhaps the most typical application scenario (motivating also our work) is route planning in road networks, where the travel-time for traversing an arc $a = uv$ (modeling a road segment) depends on the temporal traffic conditions while traversing uv, and

* This work is partially supported by the EU FP7/2007-2013, under grant agreements no. 609026 (project MOVESMART) and no. 621133 (project HoPE), and by DFG under grant agreement no. FOR-2083.

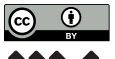

© Spyros Kontogiannis, Dorothea Wagner and Christos Zaroliagis;
licensed under Creative Commons License CC-BY

thus on the departure time from its tail u. This gives rise to *time-varying* network models and to computing min-cost (a.k.a. shortest) paths in such networks. Several variants of this problem try to model time-variation of the underlying graph structure and/or the arc-cost metric (e.g., dynamic shortest paths, parametric shortest paths, stochastic shortest paths, temporal networks, etc). In this work, we consider the case in which the cost variation of an arc a is determined by a *function* $D[a]$, which is considered to be a continuous, piecewise linear (pwl)[1] and periodic function of the time at which the resource is actually being used [7, 8, 12, 19]. In case of providing route plans in time-dependent road networks, arc-costs are considered as *arc-travel-time* functions, while time-dependent shortest paths as *minimum-travel-time* paths. The goal is then to determine the cost (*minimum-travel-time*) of a shortest path from an origin o to a destination d, as a function of the *departure-time* t_o from o. Due to the time-dependence of the arc-cost metric, the actual arc-cost value of an arc $a = uv$ is unknown until the exact time $t_u \geq t_o$ at which uv starts being traversed.

Problem Setting and Related Work. Two variants of the *time-dependent shortest path* problem have been considered in the literature: $TDSP(o, d, t_o)$ (resp. $TDSP(o, \star, t_o)$) focuses on the one-to-one (resp. one-to-all) determination of the *scalar cost* of a minimum-travel-time (shortest) path to d (resp. for all d), when departing from the given origin o at time t_o. $TDSP(o, d)$ (resp. $TDSP(o, \star)$) focuses on the one-to-one (resp., one-to-all) succinct representation of the time-dependent minimum-travel-time path *function*(s) $D[o, d]$ from o to d (resp. towards all reachable d), and all departure-times from o.

$TDSP(o, d, t_o)$ has been studied as early as [5]. The first work on $TDSP(o, d, t_o)$ for continuous time-axis was [11] where it was proved that, if *waiting-at-nodes* is allowed unconditionally, then $TDSP(o, d, t_o)$ is solvable in quasilinear time via a time-dependent variant of Dijkstra's algorithm (we call it `TDD`), which relaxes arcs by computing the arc costs "on the fly", upon settling their tails. A more complete treatment for the continuous case, considering various limitations in the waiting-times at the nodes of the network, was provided in [13]; an algorithm was also given for $TDSP(o, d, t_o)$, whose complexity cannot be bounded by a function of the network topology. An excellent overview of the problem is provided in [19]. Among other results, it was proved that for *affine arc-cost* functions possessing the FIFO property (according to which all the arc-cost functions have slopes at least -1), in addition to `TDD`, a time-dependent variant of the label-correcting Bellman-Ford algorithm works. Moreover, if waiting-at-nodes is *forbidden* and the arc-costs do not preserve the FIFO property, then *subpath-optimality* of shortest paths is not necessarily preserved. In such a case, many variants of the problem are also NP-hard [22]. Additionally, when shortest path costs are well defined and optimal waiting-times at nodes always exist, a non-FIFO arc with *unrestricted-waiting-at-tail* policy is equivalent to a FIFO arc in which waiting at the tail is not beneficial [19]. For these reasons, we focus here on instances for which the FIFO property holds, as indeed is the case with most of past and recent work on $TDSP(o, d, t_o)$.

The complexity of $TDSP(o, d)$ was first questioned in [6, 7] and remained open until recently, when it was proved in [12] that, in case of FIFO-abiding pwl arc-cost functions, for a *single* origin-destination pair (o, d) the space complexity for succinctly representing $D[o, d]$ is $(1 + K) \cdot n^{\Theta(\log n)}$, where n is the number of vertices and K is the total number of breakpoints (or legs) of all the arc-cost functions. Note that K can be substituted by the

[1] Major car navigator vendors provide real-time estimations of travel-time values by periodically sampling the average speed of road segments using the cars connected to the service as sampling devices. The most customary way to represent this historic traffic data, is to consider the continuous pwl interpolants of the sample points as *arc-travel-time functions* of the corresponding instance.

number K^* of *concavity-spoiling* breakpoints (at which the arc-cost slopes increase) of the arc-cost functions. Several *output-sensitive* algorithms for the exact computation of $D[o,d]$ have been presented [7, 8, 12, 19], the most efficient being the ones in [8, 12].

Due to the above mentioned hardness of $TDSP(o,d)$, and also since the time-dependent arc-costs are typically only (e.g., pwl) approximations of the actual costs, it is quite natural to seek for succinct representations of approximations to $D[o,d]$, which aim at trading-off accuracy for computational effort. Several *one-to-one* $(1+\varepsilon)$-approximation algorithms for $TDSP(o,d)$ appeared recently in the literature [8, 12, 18], the most successful being those provided in [18]. The first *one-to-all* $(1+\varepsilon)$-approximation algorithm for $TDSP(o,\star)$, called *bisection* (`BIS`), was given in [16]. It is based on bisecting the (common to all functions) axis of departure-times from o and considers slightly stricter assumptions than just the FIFO property for the arc-cost metric. `BIS` requires $\mathcal{O}\left(\frac{K^*}{\varepsilon}\cdot\log^2\left(\frac{n}{\varepsilon}\right)\right)$ calls to $TDSP(o,\star,t_o)$. Note that all one-to-one approximation algorithms for $TDSP(o,d)$ [8, 12, 18] demand, in worst-case, a comparable amount of calls to $TDSP(o,\star,t_o)$, just for one *od*-pair.

Minimum-travel-time oracles for time-dependent networks (TD-oracles henceforth) had received no attention until recently [16]. A TD-oracle is an offline-precomputed data structure that allows the evaluation of an *upper-approximation* $\overline{\Delta}[o,d](t_o)$ of $D[o,d](t_o)$ (due to the hardness of computing and storing $D[o,d]$), for *any* possible query $(o,d,t_o)\in V\times V\times\mathbb{R}_{\geq 0}$ that may appear in an online fashion. One trivial solution would be to provide a succinct representation of $\overline{\Delta}[o,d]$ for each pair $(o,d)\in V\times V$, for the sake of rapid evaluations in the future but at the expense of superquadratic space. Another trivial solution would be to execute `TDD` "on-the-fly" for each new query (o,d,t_o), at the expense of superlinear query-time. A non-trivial TD-oracle should aim to trade-off smoothly preprocessing requirements with query-response times and approximation guarantees. In particular, it should precompute a data structure in *subquadratic* time and space, and also provide a query algorithm which evaluates *efficiently* (i.e., faster than `TDD`) $\overline{\Delta}[o,d](t_o)$, where $\overline{\Delta}[o,d]$ must have a *provably* good approximation guarantee w.r.t. $D[o,d]$. Note that there exists important applied work (*speedup heuristics*) for computing time-dependent shortest paths (e.g., [4, 9, 10, 17]), which however provide mainly empirical evidence on the success of the adopted approaches.

The TD-oracles in [16] required $\mathcal{O}\big(n^{2-\beta}(K^*+1)\big)$ preprocessing space and time, for some constant $\beta\in(0,1)$, and are able to answer queries (under certain conditions) in time $\mathcal{O}\big(n^{\delta}\big)$, for some constant $\delta\in(0,1)$. When $K^*\in o(n)$, the oracles can be fine-tuned so as to assure sublinear query-times and subquadratic preprocessing requirements in n. An extensive experimental evaluation of those oracles on a real-world road network is provided in [14], demonstrating their practicality, at the expense, however, of large memory consumption due to the linear dependence of the preprocessing space on K^* which can be $\Omega(n)$.

The main challenge addressed here is to provide TD-oracles that achieve: (i) subquadratic preprocessing requirements, *independently* of K^*; and (ii) query-times sublinear, not only in the worst-case (i.e., in n), but also in the number $\Gamma[o,d](t_o)$ of settled vertices when executing `TDD`$(o,\star,t_o)$ until d is settled (*Dijkstra-Rank*).

Our Contributions. We address positively the aforementioned challenge by providing: *(i)* A novel and remarkably simple algorithm (`TRAP`) (cf. Section 3) for constructing one-to-many $(1+\varepsilon)$-upper-approximations $\overline{\Delta}[o,d]$ (*summaries*) of minimum-travel-time functions $D[o,d]$, for all "sufficiently distant" destinations d from the origin o. `TRAP` requires $o(n)$ calls to $TDSP(o,\star,t_o)$, which is *independent* of the degree of concavity K^*. Its novelty is that it does not demand the concavity of the unknown function to approximate. *(ii)* The `TRAPONLY` and `FLAT` oracles (cf. Section 4) which exploit `TRAP` and `BIS` to construct minimum-travel-time summaries from randomly selected landmarks to *all* reachable destinations. The pre-

	preprocessing space/time	query time	recursion budget (depth) r
[16]	$K^* \cdot n^{2-\beta+o(1)}$	$n^{\delta+o(1)}$	$r \in \mathcal{O}(1)$
TRAPONLY	$n^{2-\beta+o(1)}$	$n^{\delta+o(1)}$	$r \approx \frac{\delta}{a} - 1$
FLAT	$n^{2-\beta+o(1)}$	$n^{\delta+o(1)}$	$r \approx \frac{2\delta}{a} - 1$
HORN	$n^{2-\beta+o(1)}$	$\approx \Gamma[o,d](t_o)^{\delta+o(1)}$	$r \approx \frac{2\delta}{a} - 1$

Table 1 Achievements of oracles for TD-instances with period $T = n^a$, for constant $a \in (0,1)$. The stretch of all query algorithms is $1 + \varepsilon \cdot \frac{(\varepsilon/\psi)^{r+1}}{(\varepsilon/\psi)^{r+1}-1}$. For all oracles, except for the first, we assume that $\beta \downarrow 0$.

processed data structures require subquadratic space and time, independently of K^*. FLAT uses the query algorithms in [16]. TRAPONLY needs to extend them in order to recover missing summaries for local neighborhoods around a landmark. In both cases sublinear query-times are achieved. *(iii)* The HORN oracle (cf. Section 5) which organizes a hierarchy of landmarks, from many local landmarks possessing summaries only for small neighborhoods of destinations around them, up to a few global landmarks possessing summaries for all reachable destinations. HORN's preprocessing requirements are again subquadratic. We then devise and analyze a novel query algorithm (HQA) which exploits this hierarchy, with query-time *sublinear* in the Dijkstra-Rank of the query at hand.

Except for the choice of landmarks, our algorithms are deterministic. A recent experimental study [15] demonstrates the excellent performance of our oracles in practice, achieving considerable memory savings and query times about three orders of magnitude faster than TDD, and more than 70% faster than those in [14]. Table 1 summarizes the achievements of the TD-oracles presented here and their comparison with the oracles in [16].

2 Preliminaries

Notation and Terminology. For any integer $k \geq 1$, let $[k] = \{1, 2, \ldots, k\}$. A *time-dependent network instance* (TD-instance henceforth) consists of a directed graph $G = (V, A)$ with $|V| = n$ vertices and $|A| = m \in \mathcal{O}(n)$ arcs, where each arc $a \in A$ is accompanied with a continuous, pwl *arc-cost* function $D[a] : \mathbb{R}_{\geq 0} \mapsto \mathbb{R}_{\geq 0}$. We assume that all these functions are periodic with period $T > 0$ and are defined as follows: $\forall k \in \mathbb{N}, \forall t \in [0,T)$, $D[a](kT + t) = d[a](t)$, where $d[a] : [0,T) \to (0, M_a]$ is such that $\lim_{t \uparrow T} d[a](t) = d[a](0)$, for some fixed integer M_a denoting the maximum possible cost ever seen for arc a. Let also $M = \max_{a \in A} M_a$ denote the maximum arc-cost ever seen in the entire network. Since $D[a]$ is periodic, continuous and pwl function, it can be represented succinctly by a sequence of K_a breakpoints (i.e., pairs of departure-times and arc-cost values) defining $d[a]$. $K = \sum_{a \in A} K_a$ is the number of breakpoints representing all arc-cost functions, $K_{\max} = \max_{a \in A} K_a$, and K^* is the number of *concavity-spoiling* breakpoints (the ones at which the arc-cost function slopes increase). Clearly, $K^* \leq K$, and $K^* = 0$ for *concave* arc-cost functions.

To ease the exposition and also for the sake of compliance with terminology in previous works (inspired by the primary application scenario of route planning in time-dependent road networks), we consider arc-costs as *arc-travel-times* and time-dependent shortest paths as *minimum-travel-time* paths. This terminology facilitates the following definitions. The *arc-arrival-time* function of $a \in A$ is $Arr[a](t) = t + D[a](t)$, $\forall t \in [0, \infty)$. The *path-arrival-time* function of a path $p = \langle a_1, \ldots, a_k \rangle$ in G (represented as a sequence of arcs) is the composition $Arr[p](t) = Arr[a_k](Arr[a_{k-1}](\cdots(Arr[a_1](t))\cdots))$ of the arc-arrival-time functions for the constituent arcs. The *path-travel-time* function is then $D[p](t) = Arr[p](t) - t$.

For any $(o,d) \in V \times V$, $\mathcal{P}_{o,d}$ denotes the set of *od*-paths. For any $p \in \mathcal{P}_{o,x}$ and $q \in \mathcal{P}_{x,d}$, $s = p \bullet q \in \mathcal{P}_{o,d}$ is the concatenation of p and q at x. The *earliest-arrival-time* function is $Arr[o,d](t_o) = \min_{p \in \mathcal{P}_{o,d}} \{Arr[p](t_o)\}$, $\forall t_o \geq 0$, while the *minimum-travel-time* function is defined as $D[o,d](t_o) = \min_{p \in \mathcal{P}_{o,d}} \{D[p](t_o)\} = Arr[o,d](t_o) - t_o$. For a given query (o,d,t_o), $SP[o,d](t_o) = \{p \in P_{o,d} : Arr[p](t_o) = Arr[o,d](t_o)\}$ is the set of earliest-arrival-time (equivalently, minimum-travel-time) paths. $ASP[o,d](t_o)$ is the set of *od*-paths whose travel-time values are $(1+\varepsilon)$-approximations of the minimum-travel-time among all *od*-paths.

When we say that we *"grow a* `TDD` *ball from* (o,t_o)*"*, we refer to the execution of `TDD` from $o \in V$ at departure time $t_o \in [0,T)$ for solving $TDSP(o,\star,t_o)$ (resp. $TDSP(o,d,t_o)$, for a specific destination d). Such a call, which we denote as $\texttt{TDD}(o,\star,t_o)$ (resp. $\texttt{TDD}(o,d,t_o)$), takes time $\mathcal{O}(m + n\log(n)[1 + \log\log(1 + K_{\max})]) = \mathcal{O}(n\log(n)\log\log(K_{\max})])$, using predecessor search for evaluating continuous pwl functions (cf. [16]). The *Dijkstra-Rank* $\Gamma[o,d](t_o)$ is the number of settled vertices up to d, when executing $\texttt{TDD}(o,d,t_o)$.

$\forall a = uv \in A$ and $[t_s,t_f) \subseteq [0,T)$, we define upper- and lower-bounding travel-time metrics: the *minimally-congested* travel-time $\underline{D}[uv](t_s,t_f) := \min_{t_u \in [t_s,t_f)}\{D[uv](t_u)\}$ and the *maximally-congested* travel-time $\overline{D}[uv](t_s,t_f) := \max_{t_u \in [t_s,t_f)}\{D[uv](t_u)\}$. If $[t_s,t_f) = [0,T)$, we refer to the static *free-flow* and *full-congestion* metrics $\underline{D}, \overline{D} : A \to [1,M]$, respectively. Each arc $a \in A$ is also equipped with scalars $\underline{D}[a]$ and $\overline{D}[a]$ in these static metrics. For any arc-cost metric D, $diam(G,D)$ is the diameter (largest possible vertex-to-vertex distance) of the graph. For example, $diam(G,\underline{D})$ is the free-flow diameter of G.

In our TD-instance, we can guarantee that $T \geq diam(G,\underline{D})$. If this is not the case, we can take the minimum number c of consecutive copies of each $d[a]$ as a single arc-travel-time function $d'[a] : [0,cT) \mapsto \mathbb{R}_{>0}$ and $D'[a](t+kT') = d'[a](t)$, $\forall t \in [0,T')$ such that $T' = cT \geq diam(G,\underline{D'})$. In addition, we can also guarantee that $T = n^\alpha$ for a small *constant* $\alpha \in (0,1)$ of our control. If $T \neq n^\alpha$, we scale the travel-time metric by setting $D'' = \frac{n^\alpha}{T} \cdot D$ (e.g., we change the unit by which we measure time, from milliseconds to seconds, or even minutes) and use the period $T'' = n^\alpha$, without affecting the structure of the instance at all. From now on we consider w.l.o.g. TD-instances with $T = n^\alpha \geq diam(G,\underline{D})$.

For any $v \in V$, departure-time $t_v \in \mathbb{R}_{\geq 0}$, integer $F \in [n]$ and $R > 0$, $B[v;F](t_v)$ ($B[v;R](t_v)$) is a ball of size F (of radius R) grown by `TDD` from (v,t_v), in the time-dependent metric. Analogously, $\underline{B}[v;F]$ ($\underline{B}[v;R]$) and $\overline{B}[v;F]$ ($\overline{B}[v;R]$) are, respectively, the size-F (radius-R) balls from v in the free-flow and fully-congested travel-time metrics.

A pair of continuous, pwl, periodic functions $\overline{\Delta}[o,d]$ and $\underline{\Delta}[o,d]$), with a (hopefully) small number of breakpoints, are $(1+\varepsilon)$-*upper-approximation* and $(1+\varepsilon)$-*lower-approximation* of $D[o,d]$, if $\forall t_o \geq 0$, $\frac{D[o,d](t_o)}{1+\varepsilon} \leq \underline{\Delta}[o,d](t_o) \leq D[o,d](t_o) \leq \overline{\Delta}[o,d](t_o) \leq (1+\varepsilon) \cdot D[o,d](t_o)$.

Assumptions on the time-dependent arc-cost metric. The directedness and time-dependence in the underlying network imply an asymmetric arc-cost metric that also evolves with time. To achieve a smooth transition from static and undirected graphs towards time-dependent and directed graphs, we need a quantification of the degrees of asymmetry and evolution of our metric over time. These are captured via a set of parameters depicting the steepness of the minimum-travel-time functions, the ratio of minimum-travel-times in opposite directions, and the relation between graph expansion and travel-times. We make some assumptions on the values of these parameters which seem quite natural for our main application scenario (route planning in road networks), and were verified by an experimental analysis (cf. Appendix B). Here we only present a qualitative interpretation of them. It is noted that Assumptions 2.1 and 2.2 were exploited also in the analyses in [16].

▶ Assumption 2.1 (Bounded Travel-Time Slopes). All the minimum-travel-time slopes are bounded in a given interval $[-\Lambda_{\min}, \Lambda_{\max}]$, for given constants $\Lambda_{\min} \in [0,1)$ and $\Lambda_{\max} \geq 0$.

▶ Assumption 2.2 (Bounded Opposite Trips). The ratio of minimum-travel-times in opposite directions between two vertices, for any specific departure-time but not necessarily via the same path, is upper bounded by a given constant $\zeta \geq 1$.

▶ Assumption 2.3 (Growth of Free-Flow Dijkstra Balls). $\forall F \in [n]$, the free-flow ball $\underline{B}[v;F]$ blows-up by at most a polylogarithmic factor, when expanding its (free-flow) radius up to the value of the full-congestion radius within $\underline{B}[v;F]$.

Finally, we need to quantify the correlation between the arc-cost metric and the Dijkstra-Rank metric induced by it. For this reason, inspired by the notion of the doubling dimension (e.g., [2] and references therein), we consider some *scalar* $\lambda \geq 1$ and functions $f, g : \mathbb{N} \mapsto [1, \infty)$, such that the following hold: $\forall (o, d, t_o) \in V \times V \times [0, T)$, (i) $\Gamma[o,d](t_o) \leq f(n) \cdot (D[o,d](t_o))^{\lambda}$, and (ii) $D[o,d](t_o) \leq g(n) \cdot (\Gamma[o,d](t_o))^{1/\lambda}$. This property trivially holds, e.g., for $\lambda = 1$, $f(n) = n$, and $g(n) = \max_{a \in A} \{\overline{D}[a]\}$. Of course, our interest is for the smallest possible values of λ and at the same time the slowest-growing functions $f(n), g(n)$. Our last assumption quantifies the boundedness of this correlation by restricting λ, $f(n)$ and $g(n)$.

▶ Assumption 2.4. There exist $\lambda \in o\left(\frac{\log(n)}{\log\log(n)}\right)$ and $f(n), g(n) \in \text{polylog}(n)$ s.t. the following hold: (i) $\Gamma[o,d](t_o) \leq f(n) \cdot (D[o,d](t_o))^{\lambda}$, and (ii) $D[o,d](t_o) \leq g(n) \cdot (\Gamma[o,d](t_o))^{1/\lambda}$. Analogous inequalities hold for the free-flow and the full-congestion metrics $\underline{D}$ and $\overline{D}$.

Note that static oracles based on the doubling dimension (e.g., [2]) demand a *constant* value for the exponent λ of the expansion. We relax this by allowing λ being expressed as a (sufficiently slowly) growing function of n. We also introduce some additional slackness, by allowing some divergence from the corresponding powers by polylogarithmic factors.

In the rest of the paper we consider sparse ($m \in \mathcal{O}(n)$) TD-instances, compliant with Assumptions 2.1, 2.2, 2.3, and 2.4. For convenience, the notation used throughout the paper is summarized in Appendix A. A review of the oracles in [16] is presented in Appendix C.

3 The TRAP approximation method

We now introduce the *trapezoidal* (TRAP) algorithm, a novel algorithm for computing one-to-many $(1+\varepsilon)$-upper-approximations $\overline{\Delta}[\ell, v] : [0, T) \mapsto \mathbb{R}_{>0}$ of min-cost functions $D[\ell, v]$, from a vertex ℓ towards all sufficiently distant destinations (typically, ℓ will be a landmark). TRAP is remarkably simple and works as follows. First, $[0, T)$ is split into a number of $\lceil \frac{T}{\tau} \rceil$ consecutive length-τ subintervals, where τ is a tuning parameter to be fixed later. Then, for each such interval $[t_s, t_f = t_s + \tau) \subseteq [0, T)$, a $(1+\varepsilon)$-upper-approximation of the projection $D[\ell, v] : [t_s, t_f) \mapsto \mathbb{R}_{>0}$ is computed. Finally, the concatenation of all these $(1+\varepsilon)$-upper-approximations per subinterval constitutes the requested $(1+\varepsilon)$-upper-approximation $\overline{\Delta}[\ell, v]$ of $D[o, d] : [0, T) \mapsto \mathbb{R}_{>0}$. Note that, contrary to the BIS approximation algorithm [16], no assumption is made on the shapes of the min-cost functions to approximate within each subinterval; in particular, no assumption is made on them being concave. TRAP only exploits the fact that τ is small, along with Assumption 2.1 on the boundedness of travel-time slopes.

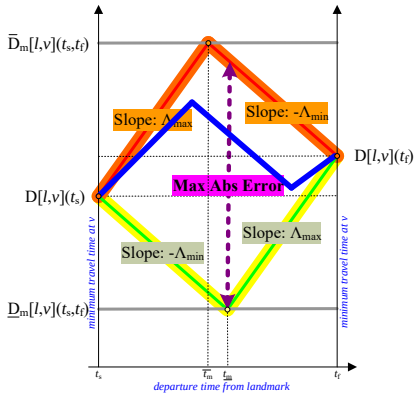

Figure 1 The upper-approximation $\overline{\delta}_k[\ell, v]$ (thick orange, upper pwl line), and lower-approximation $\underline{\delta}_k[\ell, v]$ (thick green, lower pwl line), of the unknown function $D[\ell, v]$ (blue pwl line) within the interval $I_k = [t_s = (k-1)\tau, t_f = k\tau)$.

We now describe the upper- and lower-approximations of $D[o,d]$ that we construct in a subinterval $I_k = [t_s = (k-1)\tau, t_f = k\tau) \subset [0,T)$, $k \in \left[\left\lceil \frac{T}{\tau} \right\rceil\right]$, from a vertex $\ell \in V$ towards some destination $v \in V$. The quality of the upper-approximation depends on the value of τ and the delay values at the endpoints of I_k, as we shall explain shortly. TRAP computes the following two functions of $D[\ell,v]$ (cf. Fig. 1): $\forall t \in I_k$, $\overline{\delta}_k[\ell,v](t) = \min\left\{ D[\ell,v](t_f) + \Lambda_{\min} t_f - \Lambda_{\min} t \ , \ D[\ell,v](t_s) - \Lambda_{\max} t_s + \Lambda_{\max} t \right\}$ and $\underline{\delta}_k[\ell,v](t) = \max\left\{ D[\ell,v](t_f) - \Lambda_{\max} t_f + \Lambda_{\max} t \ , \ D[\ell,v](t_s) + \Lambda_{\min} t_s - \Lambda_{\min} t \right\}$ and considers them as the upper- and lower-approximating functions of $D[\ell,v]$ within I_k. The correctness of this choice is proved in the next lemma, which follows by Assumption 2.1.

▶ **Lemma 1.** *$\overline{\delta}_k[\ell,v](t)$ and $\underline{\delta}_k[\ell,v](t)$ upper- and lower-approximate $D[\ell,v](t)$ within I_k.*

Let $(\underline{t}_m, \underline{D}_m)$ and $(\overline{t}_m, \overline{D}_m)$ be the intersections of the legs in the definition of $\underline{\delta}_k[\ell,v]$ and $\overline{\delta}_k[\ell,v]$, respectively. The *maximum additive error* $MAE(I_k)$ for $\overline{\delta}_k[\ell,v]$ in I_k (i.e., the length of the purple dashed line in Fig. 1) is $MAE(I_k) := \max_{t \in I_k} \left\{ \overline{\delta}_k[\ell,v](t) - \underline{\delta}_k[\ell,v](t) \right\} = \overline{\delta}_k[\ell,v](\underline{t}_m) - \underline{\delta}_k[\ell,v](\underline{t}_m)$. The following lemma proves that, for τ sufficiently small, $MAE(I_k)$ cannot be large. It also provides a *sufficient condition* for the value of τ so that $\overline{\delta}_k[\ell,v]$ is indeed a $(1+\varepsilon)$-upper-approximation of $D[\ell,v]$ within I_k.

▶ **Lemma 2.** *$\forall (\ell,v) \in L \times V$, $\forall k \in \left[\left\lceil \frac{T}{\tau} \right\rceil\right]$ and $I_k = [(k-1)\tau, k\tau)$, the following hold: (1) $MAE[\ell,v](I_k) \leq \Lambda_{\max} \cdot \tau$; (2) $\overline{\delta}_k[\ell,v]$ is a $(1+\varepsilon)$-upper-approximation of $D[\ell,v]$ within I_k, if $\left[\ D[\ell,v](t_s) \geq \left(\Lambda_{\min} + \frac{\Lambda_{\max}}{\varepsilon}\right) \cdot \tau \ \right] \vee \left[\ D[\ell,v](t_f) \geq \left(1 + \frac{1}{\varepsilon}\right) \Lambda_{\max} \cdot \tau \ \right]$*

For given $\tau > 0$ and $\ell \in L$, the set of *faraway destinations* from ℓ is $V[\ell](\tau) = \{v \in V : \tau[\ell,v] > \tau\}$. $\tau[\ell,v] = \frac{\underline{D}[\ell,v]}{(1+1/\varepsilon)\Lambda_{\max}}$ is a sufficient τ-value for $\overline{\delta}_k[\ell,v]$ being $(1+\epsilon)$-upper-approximation of $D[\ell,v]$ within $I_k = [(k-1)\tau[\ell,v], k\tau[\ell,v])$ (cf. Lemma 2). The next theorem proves that TRAP provides a $(1+\varepsilon)$-upper-approximation $\overline{\Delta}[\ell,v]$ for all faraway destinations from ℓ, and also estimates the preprocessing requirements of the algorithm.

▶ **Theorem 3.** *Fix $\ell \in L$, $F > f(n)$, and $\tau \in (0,T)$ s.t. $|V[\ell](\tau)| = n - F$. Let $\tau^* = \min_{v \in V[\ell](\tau)} \left\{ \frac{\underline{D}[\ell,v]}{(1+1/\varepsilon)\Lambda_{\max}} \right\}$. $\forall v \in V[\ell](\tau)$, $\overline{\Delta}[\ell,v]$ is the concatenation of all the upper-approximating functions $\overline{\delta}_k[\ell,v]$ that* TRAP *returns per subinterval $I_k = [\ t_{s_k} = (k-1)\tau^* \ , \ t_{f_k} = \min\{k\tau^*, T\} \ ) : k \in \left[\left\lceil \frac{T}{\tau^*} \right\rceil\right]$. Then, $\forall v \in V[\ell](\tau)$, $\overline{\Delta}[\ell,v]$ is a $(1+\varepsilon)$-upper-approximation of $D[\ell,v]$ in $[0,T)$, requiring at most $2\left\lceil \frac{T}{\tau^*} \right\rceil$ breakpoints. The number of calls to $TDSP(\ell,\star,t)$ for their construction is $\left\lceil \frac{T}{\tau^*} \right\rceil \leq 1 + \frac{T(1+1/\varepsilon)\Lambda_{\max}}{\min_{v \in V[\ell](\tau)}\{\underline{D}[\ell,v]\}} \in \mathcal{O}(n^a)$.*

Proof of Theorem 3. τ^* is the appropriate length for the subintervals which assures that TRAP returns $(1+\varepsilon)$-upper-approximations for all faraway destinations from ℓ. By definition it holds that $\tau^* \geq \tau$. Since $F > f(n)$, it holds that TRAP does not consider destinations at free-flow distance less than 1. To see this, fix any $v \in V$ s.t. $\underline{D}[\ell,v] \leq 1$. By Assumption 2.4, $\underline{\Gamma}[\ell,v] \leq f(n) \cdot \underline{D}[\ell,v]^\lambda \leq f(n) < F$. Thus, we can be sure that $v \notin V[\ell](\tau)$. Since $T = n^a$, we conclude that $\frac{T}{\tau^*} = \frac{T(1+1/\varepsilon)\Lambda_{\max}}{\min_{v \in V[\ell](\tau)} \underline{D}[\ell,v]} \in \mathcal{O}(n^a)$. We proceed now with the analysis of TRAP. $[0,T)$ is split into $\left\lceil \frac{T}{\tau^*} \right\rceil$ consecutive length-τ^* subintervals. Lemma 1 assures that for each $I_k = [k\tau^*, (k+1)\tau^*)$ an upper-approximating function $\overline{\delta}_k[\ell,v]$ of $D[\ell,v]$ is determined, for each $v \in V[\ell](\tau)$. The concatenation of all these functions constitutes the upper-approximating function $\overline{\Delta}[\ell,v]$ for $D[\ell,v]$ within $[0,T)$. Since $\tau[\ell,v] \geq \tau^* \Rightarrow \underline{D}[\ell,v] \geq \left(1 + \frac{1}{\varepsilon}\right) \Lambda_{\max} \tau^*$, we deduce (cf. Lemma 2) that, for all $v \in V[\ell](\tau)$, the produced upper-approximations within the consecutive length-τ^* intervals are $(1+\varepsilon)$-approximations of $D[\ell,v]$. TRAP preprocesses $\ell \in L$ (concurrently for all $v \in V[\ell](\tau)$) by making $\left\lceil \frac{T}{\tau^*} \right\rceil \in \mathcal{O}(n^a)$ calls to $TDSP(\ell,\star,t)$, to sample the endpoints of all the $\left\lceil \frac{T}{\tau^*} \right\rceil$ length-τ^* subintervals. For storing $\overline{\Delta}[\ell,v]$, it needs at most $2\left\lceil \frac{T}{\tau^*} \right\rceil$ breakpoints (there is at most one intermediate breakpoint $(\overline{t}_m, \overline{D}_m)$ per subinterval). ◀

4 Oracles with fully-informed landmarks

In this section we describe two novel oracles, with landmarks possessing summaries for all reachable destinations, excluding possibly a small neighborhood around them. We start with a random landmark set $L \subset_{\mathbf{uar}(\rho)} V$, i.e., we decide independently and uniformly at random whether each vertex is a landmark, with probability $\rho = n^{-\omega}$ for a constant $\omega \in (0,1)$. We consider as *faraway vertices* from $\ell \in L$, all the vertices at free-flow distance at most $\underline{R} = T^{\theta}$ from it, for a constant $\theta \in (0,1)$ to be determined later. $F = \max_{\ell \in L} \{|\underline{B}[\ell; \underline{R}]|\}$ is the maximum number of faraway vertices from a landmark. The next lemma shows that the main parameters we should consider w.r.t. a TD-instance are λ (cf. Assumption 2.4) and $a \in (0,1)$ s.t. $T = n^a$. All the other parameters essentially adjust their values to them.

▶ **Lemma 4.** *For $\nu \in (0,1)$ s.t. $T = diam(G, \underline{D})^{1/\nu}$, $\theta \in (0,1)$ s.t. $\frac{\nu}{\theta} \in \mathcal{O}(1)$ and λ, f, g defined as in Assumption 2.4, the following hold: (i) $\frac{1}{\lambda\nu} = \alpha \pm o(1)$, and (ii) $F \in n^{[1\pm o(1)]\theta/\nu}$.*

The* TRAPONLY *oracle. A first attempt towards avoiding the dependency of the preprocessing requirements on K^*, is to develop an oracle whose preprocessing is based solely on TRAP. The preprocessing of this oracle (we call it TRAPONLY) first splits the entire period $[0,T)$ into consecutive subintervals of length $\tau = \frac{\underline{R}}{(1+1/\varepsilon)\Lambda_{\max}} > 0$ each. It then calls TRAP for each landmark $\ell \in L$, which guarantees $(1+\varepsilon)$-upper-approximations for all the faraway destinations $v \in V[\ell](\tau)$ (cf. Theorem 3). As for the faraway destinations from ℓ, we let their distances from ℓ be computed by the query algorithm that we are using (by growing a TDD ball from ℓ). In particular, the query algorithm of TRAPONLY is an appropriate variant of RQA, we call it RQA$^+$, which additionally grows a small TDD ball of size $F\,\mathrm{polylog}(F)$ (cf. Assumption 2.3) in order to compute the actual travel-times towards their faraway destinations. The following theorem analyzes the performance of TRAPONLY.

▶ **Theorem 5.** *The expected time of* RQA$^+$ *and the preprocessing requirements of* TRAPONLY *are:* $\mathbb{E}\{Q_{\texttt{RQA}^+}\} \in \mathcal{O}\left(n^{\omega r + \max\{\omega, \frac{\theta}{\nu}\} + o(1)}\right)$ *and* $S_{\texttt{TRAPONLY}},\ P_{\texttt{TRAPONLY}} \in \mathcal{O}\left(n^{2+\alpha\cdot(1-\theta)-\omega+o(1)}\right)$.

Proof of Theorem 5. During the preprocessing, TRAPONLY makes $\lceil \frac{T}{\tau} \rceil \le 1 + \frac{T(1+1/\varepsilon)\Lambda_{\max}}{\underline{R}} = 1 + T^{1-\theta}(1+1/\varepsilon)\Lambda_{\max}$ calls of $\texttt{TDD}(\ell, t)$, for departure-times $t \in \{0, \tau, 2\tau, \ldots, \lceil \frac{T}{\tau} \rceil - 1\}$ and landmarks $\ell \in L$, where the equality comes from Lemma 4. Therefore, the preprocessing-time is dominated by the aggregate time for all these TDD probes. Taking into account that each TDD probe takes time $\mathcal{O}(n \log(n) \log\log(K_{\max}))$ and that $|L| = \rho n = n^{1-\omega}$ landmarks, by using Lemma 4 we get the following: $P_{\texttt{TRAPONLY}} = n^{1-\omega} \cdot n^{\frac{1-\theta}{\nu\lambda}[1+o(1)]} \cdot n \log(n) \log\log(n) \in n^{2-\omega+\frac{1-\theta}{\nu\lambda}[1+o(1)]+\frac{\log\log(n)+\log\log\log(n)}{\log(n)}} = n^{2-\omega+a\cdot(1-\theta)+o(1)}$. The calculations are analogous for the required preprocessing space: For all landmarks $\ell \in L$ and all their faraway destinations $v \in V[\ell](\tau)$, the total number of breakpoints to store is at most $S_{\texttt{TRAPONLY}} = 2\lceil \frac{T}{\tau} \rceil \rho n^2 \in n^{2-\omega+\frac{1-\theta}{\nu\lambda}[1+o(1)]+o(1)} = n^{2-\omega+a\cdot(1-\theta)+o(1)}$. As for the query-time complexity of RQA$^+$, recall that the expected number of TDD balls that it grows is $(1/\rho)^r$. Additionally, RQA$^+$ grows $(1/\rho)^r$ TDD balls from the corresponding closest landmarks. Each ball from a new center costs $\mathcal{O}((1/\rho)\log(1/\rho))$. Each ball from a landmark costs $\mathcal{O}(F\,\mathrm{polylog}(F)) \in n^{[1\pm o(1)]\theta/\nu}$. Thus, the expected query-time is upper-bounded as follows: $\mathbb{E}\{Q_{\texttt{RQA}^+}\} \in \mathcal{O}((1/\rho)^r[(1/\rho)\log(1/\rho) + F\,\mathrm{polylog}(F)]\log\log(K_{\max})) = \mathcal{O}\left(n^{\omega r + \max\{\omega, [1+o(1)]\theta/\nu\}}\right)$. ◀

The next corollaries are parameter-tuning examples showcasing the trade-offs among the sublinearity of query-time, the subquadratic preprocessing requirements and the stretch.

▶ **Corollary 6.** *For $\delta \in (\alpha, 1)$, $\beta \in (0, \alpha^2\nu]$, $\omega = \frac{\delta}{r+1}$, $\theta = \frac{\delta\nu}{r+1}$ and $r = \left\lfloor \frac{\delta\cdot(1+\alpha\nu)}{\alpha+\beta} \right\rfloor - 1$, $S_{\texttt{TRAPONLY}}, P_{\texttt{TRAPONLY}} \in n^{2-\beta+o(1)}$, $\mathbb{E}\{Q_{\texttt{RQA}^+}\} \in n^{\delta+o(1)}$ and the stretch is $1 + \varepsilon \cdot \frac{(1+\varepsilon/\psi)^{r+1}}{(1+\varepsilon/\psi)^{r+1}-1}$.*

▶ **Corollary 7.** *For any integer $k \geq 2$, let $\eta(k) = \left\lceil \frac{\log(k/(k-1))}{\log(1+\varepsilon/\psi)} \right\rceil - 1$, $\delta \in (0,1)$ and $\beta \in \left(0, \frac{\delta}{\eta(k)+2}\right)$. Then* TRAPONLY *achieves stretch $1 + k \cdot \varepsilon$ with $S_{\texttt{TRAPONLY}}, P_{\texttt{TRAPONLY}} \in n^{2-\beta+o(1)}$ and $\mathbb{E}\{Q_{\texttt{RQA}^+}\} \in n^{\delta+o(1)}$, by scaling the TD-instance so that $T = n^\alpha$, for $\alpha = \frac{\delta-[\eta(k)+2]\cdot\beta}{\eta(k)+2-\delta\nu}$.*

The* FLAT *oracle. Our second attempt, the FLAT oracle, provides preprocessed information for all reachable destinations from each landmark. In particular, it uses the query algorithm RQA [16] The preprocessing phase of FLAT is based on a proper combination of BIS and TRAP for constructing travel-time summaries. Each landmark $\ell \in L$ possesses summaries for all reachable destinations: BIS handles all the (at most $F = \max_{\ell \in L}\{|\underline{B}[\ell;\underline{R}]|\}$) nearby destinations in $\underline{B}[\ell;\underline{R}]$, whereas TRAP handles all the faraway destinations of $V \setminus \underline{B}[\ell;\underline{R}]$. The space requirements for the summaries created by TRAP are exactly the same as in TRAPONLY. As for the summaries computed by BIS, we avoid the linear dependence of BIS on K^* by assuring that F is sufficiently small (but not too small) and exploiting Assumption 2.3 which guarantees that the involved subgraph $\underline{B}'[\ell;F]$ in the preprocessing phase of BIS on behalf of ℓ has size $\mathcal{O}(F\,\text{polylog}(F))$. The next lemma shows exactly that BIS is only affected by the concavity-spoiling breakpoints of arc-travel-time functions in $\underline{B}'[\ell;F]$, rather than the entire graph.

▶ **Lemma 8.** $\forall(\ell,v) \in L \times \underline{B}[\ell;F], \forall u \in V \setminus \underline{B}'[\ell;F], \forall t \in [0,T), D[\ell,v](t) < D[\ell,u](t)$.

Proof of Lemma 8. From the definitions of the involved free-flow and full-congestion Dijkstra balls, the following holds: $D[\ell,v](t) \leq \overline{D}[\ell,v] \leq \overline{R}[\ell] < \underline{D}[\ell,u] \leq D[\ell,u](t)$. ◀

The following theorem summarizes the complexities of the FLAT oracle.

▶ **Theorem 9.** *The query-time $Q_{\texttt{RQA}}$ and the preprocessing time $P_{\texttt{FLAT}}$ and space $S_{\texttt{FLAT}}$ of* FLAT *are: $\mathbb{E}\{Q_{\texttt{RQA}}\} \in \mathcal{O}\left(n^{\omega(r+1)+o(1)}\right)$ and $P_{\texttt{FLAT}}$, $S_{\texttt{FLAT}} \in \mathcal{O}\left(n^{1-\omega+o(1)} \cdot [n^{2\theta/\nu} + n^{1+\alpha\cdot(1-\theta)}]\right)$.*

Proof of Theorem 9. BIS requires space at most $F^2\,\text{polylog}(F)$, since by Lemma 8 the involved graph only contains $F\,\text{polylog}(F)$ vertices and concavity-spoiling breakpoints at the arc-travel-time functions. For the faraway vertices of $V \setminus \underline{B}[\ell;F]$, by setting $\tau = \frac{\underline{R}}{(1+1/\varepsilon)\Lambda_{\max}}$, TRAP provides $(1+\varepsilon)$-approximate summaries, since $\forall v \in V \setminus \underline{B}[\ell;\underline{R}]$ the sufficient condition of Theorem 3 holds: $\underline{D}[\ell,v] > \underline{R} = (1+1/\varepsilon)\,\Lambda_{\max}\tau$. Thus, we conclude that $S_{\texttt{FLAT}} \in \rho n \left[F^2\,\text{polylog}(F) + \frac{T(1+1/\varepsilon)\Lambda_{\max} n}{\underline{R}}\right] \overset{/*\text{ L.4 }*/}{=} n^{1-\omega}[n^{(2\theta/\nu)\cdot[1+o(1)]} + n^{1+\alpha\cdot(1-\theta)[1+o(1)]}]$ $= n^{1-\omega+[1+o(1)]\cdot\max\{\ 2\theta/\nu\ ,\ 1+\alpha(1-\theta)\ \}+o(1)}$, since $f(n), g(n) \in \text{polylog}(n)$. ◀

The next corollaries are parameter-tuning examples to showcase the effectiveness of FLAT.

▶ **Corollary 10.** *If $\delta \in (\alpha,1)$, $\beta \in \left(0, \frac{\alpha\cdot(1+\alpha)}{2/\nu+\alpha}\right]$, $\omega = \frac{\delta}{r+1}$, $r = \left\lfloor \frac{\delta}{\alpha} \cdot \frac{2/\nu+\alpha}{(\beta/\alpha)\cdot(2/\nu+\alpha)+(2/\nu-1)} \right\rfloor - 1$ and $\theta = \frac{1+\alpha}{2/\nu+\alpha}$, then* FLAT *has $P_{\texttt{FLAT}}, S_{\texttt{FLAT}} \in n^{2-\beta+o(1)}$, $\mathbb{E}\{Q_{\texttt{RQA}}\} \in n^{\delta+o(1)}$ and stretch $1 + \varepsilon \cdot \frac{(1+\varepsilon/\psi)^{r+1}}{(1+\varepsilon/\psi)^{r+1}-1}$.*

▶ **Corollary 11.** *For integer $k \geq 2$, let $\eta(k) = \left\lceil \frac{\log(k/(k-1))}{\log(1+\varepsilon/\psi)} \right\rceil - 1$ and $\delta \in (0,1)$.* FLAT *achieves a target stretch $1 + k \cdot \varepsilon$ with preprocessing requirements $n^{2-o(1)}$ and expected query-time $n^{\delta+o(1)}$, by scaling the TD-instance so that $T = n^\alpha$ for $\alpha = \frac{2\delta}{[\eta(k)+2]\cdot(2-\nu)-\delta\nu}$, as $\beta \downarrow 0$.*

***Comparison of* TRAPONLY *and* FLAT.** Both TRAPONLY and FLAT depend on the travel-time metric, but are independent of the degree of disconcavity K^*. On one hand, TRAPONLY is a simpler oracle, at least w.r.t. its preprocessing phase. On the other hand, FLAT achieves a

better approximation for the same TD-instance and anticipations for sublinear query-time n^{δ} and subquadratic preprocessing requirements $n^{2-\beta}$. This is because, as $\beta \downarrow 0$, FLAT guarantees a recursion budget r of (roughly) $\frac{2\delta}{a} - 1$, whereas TRAPONLY achieves about half this value and r has an exponential effect on the stretch that the query algorithms achieve.

5 The HORN oracle

We now describe and analyze the *Hierarchical ORacle for time-dependent Networks* (HORN), whose query algorithm is highly competitive against TDD, not only for long-range queries (i.e., having Dijkstra-Rank proportional to the network size) but also for medium- and short-range queries, while ensuring *subquadratic* preprocessing space and time. The main idea of HORN is to preprocess: many landmarks, each possessing summaries for a few destinations around them, so that all short-range queries can be answered using only these landmarks; fewer landmarks possessing summaries for more (but still not all) destinations around them, so that medium-range queries be answered by them; and so on, up to only a few landmarks (those required by FLAT) possessing summaries for all reachable destinations. The *area of coverage* $C[\ell] \subset V$ of ℓ is the set of its nearby vertices, for which ℓ possesses summaries. ℓ is called *informed* for each $v \in C[\ell]$, and *uninformed* for each $v \in V \setminus C[\ell]$. The landmarks are organized in a hierarchy, according to the sizes of their areas of coverage. Each level L_i in the hierarchy is accompanied with a *targeted Dijkstra-Rank* $N_i \in [n]$, and the goal of HORN is to assure that L_i should suffice for RQA to successfully address queries (o, d, t_o) with $\Gamma[o, d](t_o) \leq N_i$, in time $o(N_i)$. The difficulty of this approach lies in the analysis of the query algorithm. We want to execute a variant of RQA which, based on a *minimal* subset of landmarks, would guarantee a $(1 + \sigma(r))$-approximate solution for any query (o, d, t_o) (as in FLAT), but also time-complexity sublinear in $\Gamma[o, d](t_o)$. We propose the *Hierarchical Query Algorithm* (HQA) which grows an initial ball from (o, t_o) that stops only when it settles an informed landmark ℓ w.r.t. d which is at the "right distance" from o, given the density of landmarks belonging to the same level with ℓ. HQA essentially "guesses" as appropriate level-i in the hierarchy the level that contains ℓ, and continues with the execution of RQA with landmarks having coverage at least equal to that of ℓ (cf. Fig. 2).

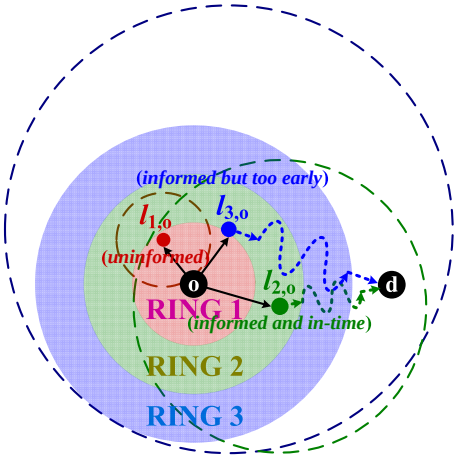

Figure 2 Demonstration of execution of HQA. Dashed circles indicate areas of coverage. Solid circular stripes indicate the rings of the corresponding levels in the hierarchy. Landmark $\ell_{1,o}$ is uninformed and $\ell_{3,o}$, although informed, comes too early. $\ell_{2,o}$ is both informed and within the ring of its own level, leading HQA to deduce that the appropriate level is $i = 2$.

***Description of* HORN.** We use the following parameters for the hierarchical construction: (i) $k \in \mathcal{O}(\log\log(n))$ determines the number of levels (minus one) comprising the hierarchy of landmarks. (ii) $\gamma > 1$ determines the actual values of the targeted Dijkstra-Ranks, one per level of the hierarchy. In particular, as γ gets closer to 1, the targeted Dijkstra-Ranks accumulate closer to small- and medium-rank queries. (iii) $\delta \in (0, 1)$ is the parameter that quantifies the sublinearity of the query algorithm (HQA), in each level of the hierarchy, compared to the targeted Dijkstra-Rank of this level. In particular, if N_i is the targeted

Dijkstra-Rank corresponding to level-i in the hierarchy, then HQA should be executed in time $\mathcal{O}\big((N_i)^{\delta}\big)$, if only the landmarks in this level (or in higher levels) are allowed to be used.

Preprocessing of HORN. $\forall i \in [k]$, we set the targeted Dijkstra-Rank for level-i to $N_i = n^{(\gamma^i-1)/\gamma^i}$. Then, we construct a randomly chosen level-i landmark set $L_i \subset_{\mathbf{uar}(\rho_i)} V$, where $\rho_i = N_i^{-\delta/(r+1)} = n^{-\delta(\gamma^i-1)/[(r+1)\gamma^i]}$. Each $\ell_i \in L_i$ acquires summaries for all (and only those) $v \in C[\ell_i]$, where $C[\ell_i]$ is the smallest *free-flow* ball centered at ℓ_i containing $c_i = N_i \cdot n^{\xi_i} = n^{(\gamma^i-1)/\gamma^i+\xi_i}$ vertices, for a sufficiently small constant $\xi_i > 0$. The summaries to the $F_i = c_i^{\chi}$ nearby vertices around ℓ_i are constructed with BIS; the summaries to the remaining $c_i - F_i$ faraway vertices of ℓ_i are constructed with TRAP, where $\chi = \frac{\theta}{\nu} = \frac{1+\alpha}{2+\alpha\nu} \in \left[\frac{1}{2}, \frac{2}{2+\nu}\right]$ is an appropriate value determined to assure the correctness of FLAT w.r.t. the level-i of the hierarchy. An ultimate level $L_{k+1} \subset_{\mathbf{uar}(\rho_{k+1})} V$ of landmarks, with $\rho_{k+1} = n^{-\frac{\delta}{r+1}}$, assures that HORN is also competitive against queries with Dijkstra-Rank greater than $n^{(\gamma^k-1)/\gamma^k}$. We choose in this case $c_{k+1} = N_{k+1} = n$, $F_{k+1} = n^{\chi}$ and $C[\ell_{k+1}] = V$, $\forall \ell_{k+1} \in L_{k+1}$.

Description of HQA. A TDD ball from (o, t_o) is grown until d is settled, or the *(ESC)*-criterion or the *(ALH)*-criterion is fulfilled (whichever occurs first):

◇ **Early Stopping Criterion (ESC):** $\ell_o \in L = \cup_{i\in[k+1]} L_i$ is settled, which is informed ($d \in C[\ell_o]$) and, for $\varphi \geq 1$, $\frac{\overline{\Delta}[\ell_o,d](t_o+D[o,\ell_o](t_o))}{D[o,\ell_o](t_o)} \geq (1+\varepsilon)\cdot\varphi\cdot(r+1)+\psi-1$.

◇ **Appropriate Level of Hierarchy (ALH):** For some level $i \in [k]$ of the hierarchy, the first landmark $\ell_{i,o} \in L_i$ is settled such that: *(i)* $d \in C[\ell_{i,o}]$ ($\ell_{i,o}$ is "informed"); and *(ii)* $\frac{N_i^{\delta/(r+1)}}{\ln(n)} \leq \Gamma[o,\ell_{i,o}](t_o) \leq \ln(n)\cdot N_i^{\delta/(r+1)}$ ($\ell_{i,o}$ is at the "right distance"). In that case, HQA concludes that i is the "appropriate level" of the hierarchy to consider. Observe that the level-$(k+1)$ landmarks are always informed. Thus, if no level-$(\leq k)$ informed landmark is discovered at the right distance, then the first level-$(k+1)$ landmark that will be found at distance larger than $\ln(n)\cdot N_k^{\delta/(r+1)}$ will be considered to be at the right distance, and then HQA concludes that the appropriate level is $k+1$.

If d is settled, an exact solution is returned. If (ESC) causes termination of HQA, the value $D[o,\ell_o](t_o) + \overline{\Delta}[\ell_o,d](t_o + D[o,\ell_o](t_o))$ is reported. Otherwise, HQA stops the initial ball due to the (ALH)-criterion, considering $i \in [k+1]$ as the appropriate level, and then continues executing the variant of RQA, call it $\texttt{RQA}_i$, which uses as its landmark set $M_i = \cup_{j=i}^{k+1} L_j$. Observe that $\texttt{RQA}_i$ may fail constructing approximate solutions via certain landmarks in M_i that it settles, since they may not be informed about d. Eventually, HQA returns the best *od*-path (w.r.t. the approximate travel-times) among the ones discovered by $\texttt{RQA}_i$ via *all* settled and informed landmarks ℓ. Theorem 12 summarizes the performance of HORN.

▶ **Theorem 12.** *Consider any TD-instance with* $\lambda \in \mathrm{o}\left(\sqrt{\frac{\log(n)}{\log\log(n)}}\right)$ *and* $g(n), f(n) \in \mathrm{polylog}(n)$ *(cf. Assumption 2.4). For* $\varphi = \frac{\varepsilon\cdot(r+1)}{\psi\cdot(1+\varepsilon/\psi)^{r+1}-1}$ *and* $k \in \mathcal{O}(\log\log(n))$*, let* $\xi_i \in \left(\left[(1+\lambda)\cdot\log\log(n) + \lambda\log\left(1+\frac{\zeta}{1-\Lambda_{\min}}\right)\right]/\log(n)\ ,\ 1-\gamma^{-i}\right)$*, for all* $i \in [k]$*. Then, for any query* (o,d,t_o) *s.t.* $N_{i^*-1} < \Gamma[o,d](t_o) \leq N_{i^*}$ *for some* $i^* \in [k+1]$*, any* $\delta \in (a,1)$*,* $\beta > 0$*, and* $r = \left\lfloor \frac{\delta}{a}\cdot\frac{(2/\nu+a)(1-\gamma)}{\beta\cdot(2/(a\nu)+1)+2/\nu-1}\right\rfloor - 1$*,* HORN *achieves* $\mathbb{E}\{Q_{\texttt{HQA}}\} \in (N_{i^*})^{\delta+\mathrm{o}(1)}$*,* $P_{\texttt{HORN}}$ *,* $S_{\texttt{HORN}} \in n^{2-\beta+\mathrm{o}(1)}$ *and stretch* $1+\varepsilon\frac{(1+\varepsilon/\psi)^{r+1}}{(1+\varepsilon/\psi)^{r+1}-1}$*, with probability at least* $1-\mathcal{O}\left(\frac{1}{n}\right)$.

References

1 R. Agarwal, P. Godfrey. Distance oracles for stretch less than 2. *SODA 2013*, pp. 526-538.

2 Y. Bartal, L.A. Gottlieb, T. Kopelowitz, M. Lewenstein, L. Roditty. Fast, precise and dynamic distance queries. *SODA 2011*, pp. 840-853.

3 H. Bast, D. Delling, A. V. Goldberg, M. Müller-Hannemann, T. Pajor, P. Sanders, D. Wagner, R. Werneck. Route planning in transportation networks. Technical Report MSR-TR-2014-4, Microsoft Research, April 2015 (`http://arxiv.org/abs/1504.05140`).

4 G. V. Batz, R. Geisberger, P. Sanders, C. Vetter. Minimum time-dependent travel times with contraction hierarchies. *ACM J. of Exp. Algorithmics*, **18**, 2013.

5 K. Cooke, E. Halsey. The shortest route through a network with time-dependent intermodal transit times. *Math. Anal. and Appl.*, **14**(3):493-498, 1966.

6 B. C. Dean. Algorithms for minimum-cost paths in time-dependent networks with waiting policies. *Networks*, **44**(1):41–46, 2004.

7 B. C. Dean. Shortest paths in FIFO time-dependent networks: Theory and algorithms. Technical Report, MIT, 2004.

8 F. Dehne, O. T. Masoud, and J. R. Sack. Shortest paths in time-dependent FIFO networks. *Algorithmica*, **62**(1-2):416–435, 2012.

9 D. Delling. Time-Dependent SHARC-Routing. *Algorithmica*, **60**(1):60–94, 2011.

10 D. Delling, D. Wagner. Time-dependent route planning. *Robust and Online Large-Scale Optimization*, **LNCS 5868**, pp. 207-230. Springer, 2009.

11 S. E. Dreyfus. An appraisal of some shortest-path algorithms. *Operations Research*, **17**(3):395–412, 1969.

12 L. Foschini, J. Hershberger, S. Suri. On the complexity of time-dependent shortest paths. *Algorithmica*, **68**(4):1075–1097, 2014.

13 J. Halpern. Shortest route with tme dependent length of edges and limited delay possibilities in nodes. *Zeitschrift für Operations Research*, **21**:117–124, 1977.

14 S. Kontogiannis, G. Michalopoulos, G. Papastavrou, A. Paraskevopoulos, D. Wagner, C. Zaroliagis. Analysis and experimental evaluation of time-dependent distance oracles. *ALENEX 2015*, pp. 147-158.

15 S. Kontogiannis, G. Michalopoulos, G. Papastavrou, A. Paraskevopoulos, D. Wagner, C. Zaroliagis. Engineering oracles for time-dependent road networks. *ALENEX 2016*, pp. 1-14.

16 S. Kontogiannis, C. Zaroliagis. Distance oracles for time-dependent networks. *Algorithmica*, Vol. **74** (2016), No. 4, pp. 1404-1434.

17 G. Nannicini, D. Delling, L. Liberti, D. Schultes. Bidirectional A* search on time-dependent road networks. *Networks*, **59**:240–251, 2012.

18 M. Omran, J. R. Sack. Improved approximation for time-dependent shortest paths. *COCOON 2014*, **LNCS 8591**, Springer, pp. 453-464.

19 A. Orda, R. Rom. Shortest-path and minimum delay algorithms in networks with time-dependent edge-length. *J. of the ACM*, **37**(3):607–625, 1990.

20 M. Patrascu, L. Roditty. Distance oracles beyond the Thorup–Zwick bound. *FOCS 2010*, pp. 815-823.

21 E. Porat, L. Roditty. Preprocess, set, query! *ESA 2011*, **LNCS 6942**, Springer, pp. 603-614.

22 H. Sherali, K. Ozbay, S. Subramanian. The time-dependent shortest pair of disjoint paths problem: Complexity, models, and algorithms. *Networks*, **31**(4):259–272, 1998.

23 C. Sommer. Shortest-path queries in static networks. *ACM Comp. Surv.*, **46**, 2014.

24 C. Sommer, E. Verbin, W. Yu. Distance oracles for sparse graphs. *FOCS 2009*, pp. 703-712.

25 M. Thorup. Compact oracles for reachability and approximate distances in planar digraphs. *J. of the ACM*, **51**(6):993–1024, 2004.

26 M. Thorup, U. Zwick. Approximate distance oracles. *J. of the ACM*, **52**(1):1–24, 2005.

27 C. Wulff-Nilsen. Approximate distance oracles with improved preprocessing time. *SODA 2012*, pp. 202-208.

28 C. Wulff-Nilsen. Approximate distance oracles with improved query time. *SODA 2013*, pp. 539–549.

A Summary of Notation

Symbol	Description
$[n]$	The set of integers $\{1, 2, \ldots, n\}$, for any $n \in \mathbb{N} - \{0\}$.
$G = (V, A)$	The graph representing the underlying structure of the road network. $n = \|V\|$ and $m = \|A\|$.
$diam(G, D)$	The diameter of G under an arc-cost metric D.
$\mathcal{P}_{o,d}$	Set of od-paths in G.
$p \bullet q$	The concatenation of the ux-path p with the xv-path q at vertex x.
$ASP[o, d](t_o)$	Set of $(1+\varepsilon)$-approximations of minimum-travel-time od-paths in G, for given departure-time $t_o \geq 0$.
$SP[o, d](t_o)$	Set of minimum-travel-time od-paths in G, for given departure-time $t_o \geq 0$.
$B[v](t_v)$	A ball growing from $(v, t_v) \in V \times [0, T)$, in the time-dependent metric, until either the destination d or the closest landmark ℓ_v from (v, t_v) is settled.
$B[v; F](t_v)$	A ball growing from $(v, t_v) \in V \times [0, T)$, in the time-dependent metric, of size $F \in \mathbb{N}$.
$\overline{B}[v; F]$ / $\underline{B}[v; F]$	A ball growing from $v \in V$, in the full-congestion / free flow metric, of (integer) size $F \in \mathbb{N}$.
$\overline{B}[v; R]$ / $\underline{B}[v; R]$	A ball growing from $v \in V$, in the full-congestion / free flow metric, of (scalar) radius $R > 0$.
$B'[v; F](t_v)$	A ball growing from $(v, t_v) \in V \times [0, T)$, in the time-dependent metric, of size F polylog(F), according to Assumption 2.3.
$d[a](t)$	The limited-window arc-travel-time function for arc $a \in A$, with departure-time $t \in [0, T)$ for some *constant* time-period $T > 0$ (e.g., a single day).
M_a	Maximum possible travel-time ever seen at arc a.
M	Maximum arc-travel-time ever seen in any arc.
$D[a](t)$	Periodic arc-travel-time function for arc $a \in A$, with domain $t \in [0, \infty)$.
$Arr[a](t)$	The arc-arrival-time function for arc $a \in A$.
$D[o, d]$	Minimum-travel-time *function*, from o to d.
$\Gamma[o, d]$	Dijkstra-Ranks *function*, from o to d.
$D_{\max}[o, d]$ / $D_{\min}[o, d]$	The maximum and minimum value of $D[o, d]$.
$\overline{D}[a]$ / $\underline{D}[a]$	Travel-times of a in *full-congestion* and *free-flow* metrics, respectively.
$\overline{\Delta}[o, d]$ / $\underline{\Delta}[o, d]$	An upper-approximating / lower-approximating function to $D[o, d]$.
$Arr[o, d]$	Earliest-arrival-time *function*, from o to d.
$t_u \quad (t_v)$	Departure-time from the tail u (arrival-time at the head v) for $uv \in A$.
$TDSP(o, d, t_o)$	The problem of finding a min-cost od-path, given a departure-time t_o.
$TDSP(o, \star, t_o)$	The problem of finding a min-cost paths tree, for given departure-time t_o.
$TDSP(o, d)$	The problem of constructing a succinct representation of min-cost od-paths *function*.
K_a	Number of breakpoints in the arc-travel-time function $D[a]$.
K	Total number of breakpoints in the arc-travel-time functions.
$K_{\max}$	The maximum number of breakpoints, among the arc-travel-time functions.
K^*	Total number of *concavity-spoiling* breakpoints (i.e., points at which the slope increases) in the arc-travel-time functions.
$\Lambda_{\max}$	Maximum slope among minimum-travel-time functions.
$\Lambda_{\min}$	Absolute value of minimum slope among minimum-travel-time functions.
ζ	ratio of minimum-travel-times in opposite directions between two vertices for any specific departure-time.
r	The recursion budget for `RQA` and `HQA`.
`BIS`	The bisection approximation method for minimum-travel-time functions.
`TRAP`	The trapezoidal approximation method for minimum-travel-time functions.
`FCA`	The *Forward Constant Approximation* query algorithm.
`RQA`	The *Recursive Query Algorithm*, based on landmarks possessing information towards all possible destinations.
`HQA`	The *Hierarchical Query Algorithm*, based on a hierarchy of landmarks.
`FLAT`	The oracle that uses landmarks possessing summaries towards all possible destinations, and the `RQA` query algorithm.
`HORN`	The oracle that uses a hierarchy of landmarks with their own subset of destination vertices, and the `HQA` query algorithm.

B Assumptions on the travel-time metric

In this section, we make a few assumptions on the kind of minimum-travel-time functions in the network. All assumptions are quite natural and justified in several application scenarios, such as the urban-traffic road networks, which have motivated this work. Technically, they allow a *smooth transition* from static metrics on undirected graphs towards time-dependent metrics on directed graphs.

The first assumption, called *Bounded Travel-Time Slopes*, asserts that the partial derivatives of the minimum-travel-time functions between any pair of origin-destination vertices are bounded in a known fixed interval:

Assumption 2.1. *There exist constants* $\Lambda_{\min} \in [0,1)$ *and* $\Lambda_{\max} \geq 0$ *s.t.:* $\forall (o,d) \in V \times V,\ \forall 0 \leq t_1 < t_2,\ \frac{D[o,d](t_1) - D[o,d](t_2)}{t_1 - t_2} \in [-\Lambda_{\min}, \Lambda_{\max}]$.

The lower-bound of -1 in the minimum-travel-time function slopes is indeed a direct consequence of the FIFO property, which is typically assumed to hold in several time-dependent networks, such as road networks. $\Lambda_{\max}$ represents the maximum possible rate of change of minimum-travel-times in the network, which only makes sense to be bounded (in particular, independent of the network size) in realistic instances such as the ones representing urban-traffic time-dependent road networks.

The second assumption, called *Bounded Opposite Trips*, asserts that for any given departure time, the minimum-travel-time from o to d is not more than a *constant* $\zeta \geq 1$ times the minimum-travel-time in the opposite direction (but not necessarily along the reverse path).

Assumption 2.2. $\exists \zeta \geq 1,\ \forall (o,d) \in V \times V,\ \forall t \in [0,T),\ D[o,d](t) \leq \zeta \cdot D[d,o](t)$.

This is quite natural in road networks, (i.e., it is most unlikely that a trip in one direction is more than 10 times longer than the trip in the opposite direction during the same time period).

A third assumption concerns the relation of the Dijkstra ranks of cocentric balls in the network, with respect to the (static) *free-flow* metric implied by the time-dependent instance at hand. Its purpose is to bridge the gap between expansion of graph distances (densities of Dijkstra balls) and travel-times in the network. It essentially asserts that, given a particular origin-vertex, if one considers a free-flow ball of a certain size (and travel-time radius) and then expands further this ball to a larger radius (equal to the full-congestion radius in the free-flow ball) then the ball size that we get changes by at most a polylogarithmic factor.

Assumption 2.3. *For any vertex* $\ell \in V$*, and a positive integer* F*, consider the (static) Dijkstra ball* $\underline{B}[\ell; F]$ *around* ℓ *under the free-flow metric. Let* $\underline{R}[\ell] = \max\{\underline{D}[\ell, v] : v \in \underline{B}[\ell; F]\}$ *and* $\overline{R}[\ell] = \max\{\overline{D}[\ell, v] : v \in \underline{B}[\ell; F]\}$ *be the largest free-flow and full-congestion travel-times from* ℓ *to any other vertex in* $\underline{B}[\ell; F]$*. Finally, let* $\underline{B}'[\ell; F] = \{v \in V : \underline{D}[\ell, v](0,T) \leq \overline{R}[\ell]\}$ *be the free-flow ball around* ℓ *with the (larger) radius* $\overline{R}[\ell]$*. Then it holds that* $|\underline{B}'[\ell; F]| \in \mathcal{O}(F \cdot \text{polylog}(F))$.

The aforementioned assumptions were verified through an experimental analysis on two real-world road networks, one concerning the urban-area of the city of Berlin and the other concerning the national road network of Germany. Our experimental analysis, presented in [15], shows that for the Berlin data set, $\Lambda_{\max} < 0.062$, $\zeta < 1.2$, and the maximum ball size expansion factor ≤ 6.7, while for the Germany data set $\Lambda_{\max} < 0.22$, $\zeta < 1.1$, and the maximum ball size expansion factor ≤ 8.3.

Finally, we need a systematic way to correlate the arc-cost metric (travel-times) with the Dijkstra-Rank metric induced by it. For this reason, inspired by the notion of the

doubling dimension (e.g., [2] and references therein), we consider some *scalar* $\lambda \geq 1$ and functions $f, g : \mathbb{N} \mapsto [1, \infty)$, such that the following hold: $\forall (o, d, t_o) \in V \times V \times [0, T)$, (i) $\Gamma[o, d](t_o) \leq f(n) \cdot (D[o, d](t_o))^{\lambda}$, and (ii) $D[o, d](t_o) \leq g(n) \cdot (\Gamma[o, d](t_o))^{1/\lambda}$. This property trivially holds, e.g., for $\lambda = 1$, $f(n) = n$, and $g(n) = diam(G, \overline{D})$. Of course, our interest is for the smallest possible values of λ and at the same time the slowest-growing functions $f(n), g(n)$. Our last assumption exactly quantifies the boundedness of this correlation by restricting λ, $f(n)$ and $g(n)$.

Assumption 2.4. *For the graph $G = (V, A)$ and the time-dependent arc-cost metric D that we consider, it holds that there exist $\lambda \in o\left(\frac{\log(n)}{\log\log(n)}\right)$* and $f(n), g(n) \in \text{polylog}(n)$ such that (i) $\Gamma[o, d](t_o) \leq f(n) \cdot (D[o, d](t_o))^{\lambda}$, and (ii) $D[o, d](t_o) \leq g(n) \cdot (\Gamma[o, d](t_o))^{1/\lambda}$.

Note that static oracles related to the notion of doubling dimension (e.g., [2]), demand a *constant* value for the exponent λ of the expansion, whereas we allow even a (not too fast) growing function of the network size n. The notion of expansion that we consider introduces some additional slackness, by allowing some divergence from the corresponding powers by polylogarithmic factors of the network size.

C Review of results in [16]

Review of the approach in [16]. The TD-oracle in [16] starts by first determining, for some $\rho \in (0, 1)$, a set L of ρn independently and uniformly at random selected *landmarks* (vertices acting as reference points). During the preprocessing phase, all $(1 + \varepsilon)$-upper-approximating functions (*travel-time summaries*) $\overline{\Delta}[\ell, v]$ are constructed from each landmark $\ell \in L$ towards every reachable destination $v \in V$, using the `BIS` approximation algorithm that keeps bisecting the common axis of departure-times from ℓ, until the desired approximation guarantee is achieved in each subinterval, for all destinations. It is proved in [16] that `BIS` requires

$$\mathcal{O}\left(\frac{K^*}{\varepsilon} \max_{d \in V}\left\{\log\left(\frac{T \cdot (\Lambda_{\max} + 1)}{\varepsilon \min_{t_o}\{D[o, d](t_o)\}}\right)\right\} \max_{d \in V}\left\{\log\left(\frac{\max_{t_o}\{D[o, d](t_o)\}}{\min_{t_o}\{D[o, d](t_o)\}}\right)\right\}\right)$$

calls to $TDSP(o, \star, t_o)$, for a given origin $o \in V$ and all reachable destinations from it. The following lemma clarifies this under the lens of our TD-instnace:

▶ **Lemma 13.** `BIS` *requires $\mathcal{O}\left(\frac{K^* \cdot \log^2(n)}{\varepsilon}\right)$ calls to $TDSP(o, \star, t_o)$ to provide all the summaries of minimum-travel-time functions from a given origin $o \in V$ towards all destinations at distance at least* 1 *from o.*

Proof of Lemma 13. The crucial observation, which is a direct consequence of Assumption 2.1, is that:

$$\overline{D}[o, d] - \underline{D}[o, d] \leq \Lambda_{\max} \cdot T \Rightarrow \frac{\overline{D}[o, d]}{\underline{D}[o, d]} \leq \frac{\Lambda_{\max} \cdot T}{\underline{D}[o, d]} + 1$$

Exploiting the facts that $T = n^a$, $\underline{D}[o, d] \geq 1$ and $\Lambda_{\max} \in \mathcal{O}(1)$, we conclude that `BIS` requires $\mathcal{O}\left(\frac{K^*}{\varepsilon} \cdot \log^2\left(\frac{n}{\varepsilon}\right)\right)$ calls to $TDSP(o, \star, t_o)$. ◀

Two query algorithms were proposed, `FCA` and `RQA`, which provide constant and $(1 + \sigma)$-approximations (for *constant* $\sigma > \varepsilon$) to minimum-travel-times, respectively.

`FCA` is a simple *sublinear*-time algorithm for evaluating $\overline{\Delta}[o, d](t_o)$, guaranteeing a constant approximation w.r.t. $D[o, d](t_o)$. In particular, it grows a `TDD` ball $B[o](t_o) =$

$\{x \in V : D[o,x](t_o) \leq D[o,\ell_o](t_o)\}$ from (o,t_o), until either d or the closest landmark $\ell_o \in \arg\min_{\ell \in L}\{D[o,\ell](t_o)\}$ is settled. FCA then returns either the exact travel-time value, or the approximate travel-time value via ℓ_o, $\overline{\Delta}[o,d](t_o) = D[o,\ell_o](t_o) + \overline{\Delta}[\ell_o,d](t_o + D[o,\ell_o](t_o))$, which is a guaranteed $(1+\varepsilon+\psi)$-approximation; ψ is a constant depending on ε, ζ and $\Lambda_{\max}$, but not on the size of the network. RQA improves the approximation guarantee provided by FCA, by exploiting carefully a number of recursive accesses to the preprocessed information, each of which produces (via calls to FCA) additional candidate *od*-paths. The tuning parameter $r \in \mathbb{N}$ – the *recursion budget* – is the depth of the produced recursion tree. RQA works as follows: As long as the destination vertex has not yet been discovered in the explored area around the origin, and there is still some remaining recursion budget, it "guesses" (by exhaustively searching for it) the next vertex w_k at the boundary of the current ball, along the (unknown) shortest *od*-path. Then, it grows a new TDD ball from the new center (w_k , $t_k = t_o + D[o,w_1](t_o) + D[w_1,w_2](t_1) + \cdots + D[w_{k-1},w_k](t_{k-1})$) until it reaches the closest landmark ℓ_k to it, at distance $R_k = D[w_k,\ell_k](t_k)$. ℓ_k offers an alternative *od*-path $SOL_k = P_{o,w_1} \bullet \cdots \bullet P_{w_{k-1},w_k} \bullet Q_k \bullet \Pi_k$ by a new application of FCA, where $P_{w_i,w_{i+1}} \in SP[w_i,w_{i+1}](t_i)$, $Q_k \in SP[w_k,\ell_k](t_k)$, and $\Pi_k \in ASP[\ell_k,d](t_k + R_k)$ is the approximate suffix subpath provided by the oracle. Observe that SOL_k uses a longer (optimal, if all centers lie on the unknown min-cost path) prefix-subpath $P_{o,w_1} \bullet \cdots \bullet P_{w_{k-1},w_k}$ which is then completed with a shorter approximate suffix-subpath $Q_k \bullet \Pi_k$. It is proved in [16] that the minimum-travel-time over all the discovered approximate *od*-paths discovered by RQA, is a $(1+\sigma)-$approximation of $D[o,d](t_o)$, for any constant $\sigma > \varepsilon$. The next theorem is a consequence of the analysis in [16]:

▶ **Theorem 14** ([16]). *If a TD-instance with $m \in \mathcal{O}(n)$ and compliant with Assumptions 2.1 and 2.2 is preprocessed using* BIS *for constructing travel-time summaries from ρn landmarks chosen uniformly-at-random, then the expected values of preprocessing space $S_{\texttt{BIS}}$ and time $P_{\texttt{BIS}}$, and query time $Q_{\texttt{RQA}}$ for* RQA*, are: $\mathbb{E}\{S_{\texttt{BIS}}\} \in \mathcal{O}\big(\rho n^2(K^*+1)\big)$, $\mathbb{E}\{P_{\texttt{BIS}}\} \in \mathcal{O}\big(\rho n^2(K^*+1)\log(n)\log\log(K_{\max})\big)$, and $\mathbb{E}\{Q_{\texttt{RQA}}\} \in \mathcal{O}\Big((1/\rho)^{r+1}\log(1/\rho)\log\log(K_{\max})\Big)$, where $r \in \mathbb{N}$ is the recursion depth in* RQA *(for $r=0$ we get* FCA*). For the approximation guarantees the following hold:* FCA *returns either an exact od-path, or an approximate od-path via a landmark ℓ_o s.t. $D[o,d](t_o) \leq R_o + \overline{\Delta}[\ell_o,d](t_o+R_o) \leq (1+\varepsilon)\cdot D[o,d](t_o) + \psi\cdot R_o \leq (1+\varepsilon+\psi)\cdot D[o,d](t_o)$, where $R_o = D[o,\ell_o](t_o)$ is the minimum-travel-time to the closest landmark, and $\psi = 1 + \Lambda_{\max}(1+\varepsilon)(1+2\zeta+\Lambda_{\max}\zeta) + (1+\varepsilon)\zeta$ is a cost-metric dependent constant.* RQA *returns, for given recursion budget $r \in \mathbb{N}$, an od-path that guarantees stretch $1+\sigma$, where $\sigma = \sigma(r) \leq \frac{\varepsilon\cdot(1+\varepsilon/\psi)^{r+1}}{(1+\varepsilon/\psi)^{r+1}-1}$.*

When $K^* \in \mathrm{o}(n)$ these TD-oracles of [16] achieve both *sublinear* query times and *subquadratic* preprocessing requirements. Unfortunately, experimental evidence [14] has demonstrated that it may be the case that $K^* \in \Omega(n)$.

D The TRAP approximation method

In this section, we provide the missing proofs of Section 3.

Proof of Lemma 1. By Assumption 2.1, for any departure-time $t \in I_k = [t_s = (k-1)\tau, t_f =$

$k\tau$) from ℓ and any destination vertex $v \in V$, the following inequalities hold:

$$-\Lambda_{\min} \leq \frac{D[\ell,v](t) - D[\ell,v](t_s)}{t - t_s} \leq \Lambda_{\max}$$
$$\Rightarrow \quad \boxed{-\Lambda_{\min}(t - t_s) + D[\ell,v](t_s) \leq D[\ell,v](t) \leq D[\ell,v](t_s) + \Lambda_{\max}(t - t_s)}$$

$$-\Lambda_{\min} \leq \frac{D[\ell,v](t_f) - D[\ell,v](t)}{t_f - t} \leq \Lambda_{\max}$$
$$\Rightarrow \quad \boxed{\Lambda_{\min}(t_f - t) + D[\ell,v](t_f) \geq D[\ell,v](t) \geq -\Lambda_{\max}(t_f - t) + D[\ell,v](t_f)}$$

Combining the two inequalities we get the following bounds: $\forall v \in V, \forall t \in I_k$:

$$\max\left\{ \begin{array}{c} -\Lambda_{\min}t + \Lambda_{\min}t_s + D[\ell,v](t_s) \\ \Lambda_{\max}t - \Lambda_{\max}t_f + D[\ell,v](t_f) \end{array} \right\} \leq D[\ell,v](t) \leq \min\left\{ \begin{array}{c} \Lambda_{\max}t - \Lambda_{\max}t_s + D[\ell,v](t_s) \\ -\Lambda_{\min}t + \Lambda_{\min}t_f + D[\ell,v](t_f) \end{array} \right\}$$

Exploiting the fact that each minimum-travel-time function from ℓ to any destination $v \in V$ and departure time from I_k respects the above mentioned upper and lower bounds, one could use a simple continuous, pwl approximation of $D[\ell,v]$ within this interval:

$$\forall t \in I_k,\ \overline{\delta}_k[\ell,v](t) = \min\left\{ \begin{array}{l} \Lambda_{\max}t + D[\ell,v](t_s) - \Lambda_{\max}t_s, \\ -\Lambda_{\min}t + D[\ell,v](t_f) + \Lambda_{\min}t_f \end{array} \right\}$$

I.e., we consider the lower-envelope of the lines passing via the point $(t_s, D[\ell,v](t_s))$ with the maximum slope $\Lambda_{\max}$, and the point $(t_f, D[\ell,v](t_f))$ with the minimum slope $-\Lambda_{\min}$. Analogously, we construct a lower-bounding approximation of $D[\ell,v]$ within I_k:

$$\forall t \in I_k,\ \underline{\delta}_k[\ell,v](t) = \max\left\{ \begin{array}{l} \Lambda_{\max}t + D[\ell,v](t_f) - \Lambda_{\max}t_f, \\ -\Lambda_{\min}t + D[\ell,v](t_s) + \Lambda_{\min}t_s \end{array} \right\}$$

Figure 1 shows the (upper and lower) approximations with respect to $D[\ell,v]$ within $[t_s, t_f)$. ◀

Proof of Lemma 2. Since $(\underline{t}_m, \underline{D}_m)$ is the intersection of two lines, it is easy to show that:

$$\underline{t}_m = \frac{D[\ell,v](t_s) - D[\ell,v](t_f)}{\Lambda_{\min} + \Lambda_{\max}} + \frac{\Lambda_{\min}t_s + \Lambda_{\max}t_f}{\Lambda_{\min} + \Lambda_{\max}}$$

$$\underline{D}_m = \frac{\Lambda_{\max}D[\ell,v](t_s) + \Lambda_{\min}D[\ell,v](t_f)}{\Lambda_{\min} + \Lambda_{\max}} - \frac{\Lambda_{\min} \cdot \Lambda_{\max}}{\Lambda_{\min} + \Lambda_{\max}} \cdot (t_f - t_s)$$

Analogously, $(\overline{t}_m, \overline{D}_m)$ is also the intersection of two lines. Therefore:

$$\overline{t}_m = \frac{D[\ell,v](t_f) - D[\ell,v](t_s)}{\Lambda_{\min} + \Lambda_{\max}} + \frac{\Lambda_{\min}t_f + \Lambda_{\max}t_s}{\Lambda_{\min} + \Lambda_{\max}}$$

$$\overline{D}_m = \frac{\Lambda_{\max}D[\ell,v](t_f) + \Lambda_{\min}D[\ell,v](t_s)}{\Lambda_{\min} + \Lambda_{\max}} + \frac{\Lambda_{\min}\Lambda_{\max}}{\Lambda_{\min} + \Lambda_{\max}}(t_f - t_s)$$

We start with the upper bound on the maximum absolute error:

$$\begin{aligned}
MAE[\ell,v](I_k) &\leq \overline{D}_m - \underline{D}_m \\
&= \frac{\Lambda_{\max} D[\ell,v](t_f) + \Lambda_{\min} D[\ell,v](t_s)}{\Lambda_{\min}+\Lambda_{\max}} + \frac{\Lambda_{\min}\Lambda_{\max}}{\Lambda_{\min}+\Lambda_{\max}}(t_f - t_s) \\
&- \frac{\Lambda_{\max} D[\ell,v](t_s) + \Lambda_{\min} D[\ell,v](t_f)}{\Lambda_{\min}+\Lambda_{\max}} + \frac{\Lambda_{\min}\Lambda_{\max}}{\Lambda_{\min}+\Lambda_{\max}}(t_f - t_s) \\
&= \frac{(\Lambda_{\max}-\Lambda_{\min})[D[\ell,v](t_f) - D[\ell,v](t_s)] + 2\Lambda_{\min}\Lambda_{max}(t_f - t_s)}{\Lambda_{\min}+\Lambda_{\max}} \\
&= \frac{(\Lambda_{\max}-\Lambda_{\min})[D[\ell,v](t_f) - D[\ell,v](t_s)]/(t_f - t_s) + 2\Lambda_{\min}\Lambda_{max}}{\Lambda_{\min}+\Lambda_{\max}}(t_f - t_s) \\
&\overset{/* \text{ As.2.1 } */}{\leq} \frac{(\Lambda_{\max}-\Lambda_{\min})\Lambda_{\max} + 2\Lambda_{\min}\Lambda_{max}}{\Lambda_{\min}+\Lambda_{\max}}(t_f - t_s) = \Lambda_{\max}\tau
\end{aligned}$$

Recall now about $\overline{\delta}_k[\ell,v]$ that: $\forall t \in I_k$,

$$\begin{aligned}
\overline{\delta}_k[\ell,v](t) &\leq \underline{\delta}_k[\ell,v](t) + MAE[\ell,v](I_k) \leq \underline{\delta}_k[\ell,v](t) + \Lambda_{\max}\tau \\
&\leq D[\ell,v](t) \cdot \left(1 + \frac{\Lambda_{\max}\tau}{\underline{\delta}_k[\ell,v](t)}\right)
\end{aligned}$$

Our goal is to assure that this last upper bound of $\overline{\delta}_k[\ell,v](t)$ is in turn upper-bounded by $(1+\varepsilon)\cdot D[\ell,v](t)$. Based on the expression of $\underline{\delta}_k[\ell,v](t)$, and exploiting also the fact that $\tau \geq \max\{t - t_s, t_f - t\}$, a *sufficient condition* for this to hold, is the following:

$$D[\ell,v](t_s) \geq \left(\Lambda_{\min} + \tfrac{\Lambda_{\max}}{\varepsilon}\right)\tau \quad \vee \quad D[\ell,v](t_f) \geq \left(\Lambda_{\max} + \tfrac{\Lambda_{\max}}{\varepsilon}\right)\tau$$

This sufficient condition is independent of the actual departure time $t \in I_k$, and only depends on the travel-time values at the endpoints t_s and t_f, and also on the length τ of the departure-times subinterval that we choose. ◀

E The FLAT and TRAPONLY oracles

We start by providing the missing proof of Lemma 4.

Proof of Lemma 4. All the properties mentioned in the statement of the Lemma, are consequences of Assumption 2.4. We also exploit the fact that the free-flow diameter $diam(G,\underline{D})$ of G corresponds to the maximum possible Dijkstra-Rank, which is equal to n, assuming that the graph is strongly connected, if we use as root of the Dijkstra tree the origin of a longest minimum-travel-time path in $(G,\underline{D})$. In particular, we proceed with the explanation of each property separately:

(i) Recall that $f(n), g(n) \in \text{polylog}(n) = n^{\mathcal{O}\left(\frac{\log\log(n)}{\log(n)}\right)}$ and $\lambda \in \text{o}\left(\frac{\log(n)}{\log\log(n)}\right)$. We start by providing the upper bound of T:

$$\begin{aligned}
& diam(G,\underline{D}) \leq g(n)\cdot n^{1/\lambda} \\
\Rightarrow \quad & T = (diam(G,\underline{D}))^{1/\nu} \leq (g(n))^{1/\nu}\cdot n^{1/(\nu\lambda)} \\
& = n^{\frac{1}{\nu\lambda}+\frac{1}{\nu}\cdot\mathcal{O}\left(\frac{\log\log(n)}{\log(n)}\right)} = n^{\frac{1}{\nu\lambda}\left[1+\lambda\cdot\mathcal{O}\left(\frac{\log\log(n)}{\log(n)}\right)\right]} = n^{\frac{1+o(1)}{\nu\lambda}}
\end{aligned}$$

As for the lower bound of T, we have:

$$\begin{aligned}
& diam(G,\underline{D}) \geq \left(\tfrac{n}{f(n)}\right)^{1/\lambda} \\
\Rightarrow \quad & T = (diam(G,\underline{D}))^{1/\nu} \geq \left(\tfrac{n}{f(n)}\right)^{1/(\nu\lambda)} = n^{\frac{1}{\nu\lambda}\cdot\left[1-\mathcal{O}\left(\frac{\log\log(n)}{\log(n)}\right)\right]} = n^{\frac{1-o(1)}{\nu\lambda}}
\end{aligned}$$

(ii) For the upper bound on F we have:

$$\begin{aligned} F &= \max_{\ell \in L}\{|B[\ell, \underline{R}]|\} \\ &\leq f(n)\underline{R}^{\lambda} = f(n)T^{\theta\lambda} \leq n^{\frac{\theta}{\nu}\cdot\left[1+\left(\lambda+\frac{\nu}{\theta}\right)\cdot\mathcal{O}\left(\frac{\log\log(n)}{\log(n)}\right)\right]} = n^{\frac{[1+o(1)]\theta}{\nu}} \end{aligned}$$

where the last step is because we consider instances with $\lambda \in o\left(\frac{\log(n)}{\log\log(n)}\right)$, and moreover due to the fact that we set our (yet unspecified) tuning parameter θ so that $\frac{\nu}{\theta} \in \mathcal{O}(1)$. For the lower bound on F we have:

$$\begin{aligned} \underline{R} &\leq g(n)\cdot F^{1/\lambda} \\ F &\geq \left(\frac{\underline{R}}{g(n)}\right)^{\lambda} = \frac{T^{\theta\lambda}}{(g(n))^{\lambda}} \geq n^{[1-o(1)]\frac{\theta}{\nu}-\lambda\mathcal{O}\left(\frac{\log\log(n)}{\log(n)}\right)} \\ &= n^{\frac{\theta}{\nu}\cdot\left[1-\left(1+\frac{\nu\lambda}{\theta}\right)\cdot\mathcal{O}\left(\frac{\log\log(n)}{\log(n)}\right)\right]} = n^{\frac{[1-o(1)]\theta}{\nu}} \end{aligned}$$

where the last equality is, again, valid since $\lambda \in o\left(\frac{\log(n)}{\log\log(n)}\right)$ and $\frac{\nu}{\theta} \in \mathcal{O}(1)$.

◀

E.1 Analysis of the TRAPONLY oracle.

Recall that the preprocessing of the TRAPONLY oracle is based solely on TRAP for computing travel-time summaries, while the query algorithm is an appropriate variant of RQA (we call it RQA$^+$) which additionally grows a small TDD ball as soon as it settles a new landmark, in order to compute "on the fly" the exact minimum-travel-times (rather than evaluating preprocessed summaries, which do not exist) towards the nearby destinations from it. Theorem 5 provides the performance of the TRAPONLY oracle.

Corollaries 6 and 7 explore the conditions under which sublinear query-time and/or subquadratic preprocessing complexities can be guaranteed. We provide here their proofs.

Proof of Corollary 6. We consider scaled TD-instances with $\frac{1}{\nu\lambda} = \alpha$ for some constant $\alpha \in (0,1)$. We start with the sublinearity of the query time. For arbitrary constant $\delta \in (0,1)$ such that $\delta = \max\left\{\omega\cdot(r+1), \omega r + \frac{\theta}{\nu}\right\}$, if we choose $\omega = \frac{\theta}{\nu}$ then we have $\delta = \omega\cdot(r+1) \Leftrightarrow \omega = \frac{\theta}{\nu} = \frac{\delta}{r+1} \Leftrightarrow \theta = \frac{\delta\nu}{r+1}$ and the *sublinearity* of the expected query time n^{δ} is guaranteed.

We continue with the demand for subquadratic preprocessing requirements. Fix now some $\beta \in (0,1)$. In order to assure preprocessing time and space (roughly) $n^{2-\beta+o(1)}$, it is necessary to assure that

$$\beta \leq \omega - \alpha\cdot(1-\theta) = \frac{\delta}{r+1} - \alpha + \frac{\alpha\delta\nu}{r+1} \Leftrightarrow \boxed{r \leq \frac{\delta\cdot(1+\alpha\nu)}{\alpha+\beta} - 1}$$

◀

▶ Remark. Note that the query-time performance of RQA$^+$ is equal to that of RQA (cf. [16]), if $\omega = \frac{\theta}{\nu}$. Moreover, $\beta \leq a^2\nu$ implies that $\frac{\delta\cdot(1+\alpha\nu)}{\alpha+\beta} \geq \frac{\delta}{\alpha}$.

Proof of Corollary 7. We know that, for those instances for which it works, TRAPONLY achieves stretch $1+\varepsilon\cdot\frac{(1+\varepsilon/\psi)^{r+1}}{(1+\varepsilon/\psi)^{r+1}-1}$. We must assure then that

$$\frac{(1+\varepsilon/\psi)^{r+1}}{(1+\varepsilon/\psi)^{r+1}-1} \leq k \Rightarrow r \geq \left\lceil \frac{\log\left(\frac{k}{k-1}\right)}{\log(1+\varepsilon/\psi)} \right\rceil - 1 = \eta(k)$$

Recall that the maximum recursion budget of TRAPONLY is $\left\lfloor \frac{\delta\cdot(1+\alpha\nu)}{\alpha+\beta} \right\rfloor - 1$ (cf Corollary 6). A sufficient condition for guaranteeing the required budget for having stretch $1+k\cdot\varepsilon$ is thus the following:

$$\frac{\delta\cdot(1+\alpha\nu)}{\alpha+\beta} - 2 \geq \eta(k) \Rightarrow \boxed{\alpha \leq \frac{\delta-[\eta(k)+2]\cdot\beta}{\eta(k)+2-\delta\nu}}$$

◀

E.2 Analysis of the FLAT oracle.

We provide in this section the proof of the corollaries that showcase appropriate parameter-tunings for achieving sublinear query time and subquadratic preprocessing requirements.

Proof of Corollary 10. Again we set $\omega = \frac{\delta}{r+1}$, in order to achieve (sublinear) query time n^{δ}. We then set $\frac{2\theta}{\nu} = 1+\alpha\cdot(1-\theta) \Leftrightarrow \theta = \frac{1+\alpha}{2/\nu+\alpha}$. Observe that, for this value of θ, it is guaranteed that $\frac{\nu}{\theta} = \frac{2+\alpha\nu}{1+\alpha} < 3$, as was assumed in the proof of Lemma 4. To guarantee subquadratic preprocessing requirements $n^{2-\beta+\alpha(1)}$, we must assure that:

$$2-\beta \geq 2-\omega+\alpha\cdot(1-\theta) \Leftrightarrow \beta \leq \omega-\alpha\cdot(1-\theta) \Leftrightarrow \beta \leq \frac{\delta}{r+1} - \alpha\cdot\frac{2/\nu-1}{2/\nu+\alpha}$$

$$\Leftrightarrow \quad \boxed{r \leq \frac{\delta}{\alpha}\cdot\frac{\frac{2}{\nu}+\alpha}{\frac{\beta}{\alpha}\left(\frac{2}{\nu}+\alpha\right)+\left(\frac{2}{\nu}-1\right)} - 1}$$

The approximation guarantee is, again, the one provided by the RQA query algorithm. ◀

Proof of Corollary 11. Recall that, in order to assure a stretch factor $1+k\cdot\varepsilon$, we must set the recursion budget $r \geq \left\lceil \frac{\log\left(\frac{k}{k-1}\right)}{\log\log(1+\varepsilon/\psi)} \right\rceil - 1 = \eta(k)$, as in the proof of Corollary 7.

From Theorem 9 we also have an upper bound on the recursion budget. Thus, as $\beta \downarrow 0$, a sufficient condition for the recursion budget, so that the required stretch is achieved is the following:

$$\eta(k)+2 \leq \frac{\delta}{\alpha}\cdot\frac{2/\nu+\alpha}{2/\nu-1} \Rightarrow \boxed{\alpha \leq \frac{2\delta}{[\eta(k)+2]\cdot(2-\nu)-\delta\nu}}$$

It is now straightforward that the expected query-time is indeed $n^{\delta+\alpha(1)}$, whereas the preprocessing requirements are $n^{2-\alpha(1)}$ since we consider a very small value for β. ◀

F Analysis of the HORN oracle

The construction of the travel-time summaries for HORN is based on the FLAT (BIS + TRAP) preprocessing scenario. The queries are served by the HQA query algorithm. The oracle exploits two fundamental properties:

(i) the approximation guarantee of a path via some landmark ℓ strongly depends on the relative distance of the landmark from the origin o, compared to the distance of the destination d from o;

(ii) given that the expected distance of a level-i from the origin is roughly $\frac{1}{\rho_i}$, it is rather unlikely that the first level-i landmark will appear too early or too late (i.e., outside a sufficiently wide ring-stripe around (o, t_o)).

Property (i) is exploited by the (ESC) criterion in order to handle the exceptional case where a higher-level landmark (which also happens to be informed) appears before the first informed landmark from the appropriate level. Property (ii) is actually an event that holds with high probability, as is shown in the detailed analysis of the oracle, and is exploited by the (ALH) criterion. Therefore, the event that an informed landmark appears which is also at the right distance, whereas the previously discovered landmarks (most likely of smaller levels) were uninformed, reveals an asymptotic bound for the unknown Dijkstra-Rank $\Gamma[o,d](t_o)$ of the destination.

We now provide a sequence of lemmata which will eventually be used in the proof our main technical result, concerning the complexities of HORN mentioned in Theorem 12 (cf. Section 5).

We start with an upper bound on the free-flow distance of a discovered landmark ℓ_o from d. At this point we do not require that $d \in C[\ell_o]$.

▶ **Lemma 15.** *Let $\ell_o \in L$ be a landmark discovered by* HQA*. Then it holds that $\underline{D}[\ell_o, d] \leq \frac{\zeta}{1-\Lambda_{\min}} \cdot D[o,\ell_o](t_o) + D[o,d](t_o)$.*

Proof of Lemma 15. By Assumption 2.2 we know that:

$$D[\ell_o, o](t_o) \leq \zeta \cdot D[o,\ell_o](t_o) \tag{1}$$

By Assumption 2.1 we also know that:

$$\begin{array}{rrl} \forall x > 0, & -\Lambda_{\min} \cdot x \leq & D[\ell_o,o](t_o) - D[\ell_o,o](t_o - x) \\ \Rightarrow & D[\ell_o,o](t_o - x) \leq & D[\ell_o,o](t_o) + \Lambda_{\min} \cdot x \end{array} \tag{2}$$

We look for a particular departure-time $t_o - x_o$, and the corresponding minimum-travel-time $D[\ell_o, o](t_o - x_o)$, so as to be at the origin o exactly at time t_o. That is:

$$\begin{array}{rrcl} & t_o &=& t_o - x_o + D[\ell_o,o](t_o - x_o) \\ \Rightarrow & x_o &=& D[\ell_o,o](t_o - x_o) \overset{/*\ (2)\ */}{\leq} D[\ell_o,o](t_o) + \Lambda_{\min} \cdot x_o \\ \Rightarrow & x_o &=& \leq \frac{D[\ell_o,o](t_o)}{1-\Lambda_{\min}} \overset{/*\ (1)\ */}{\leq} \frac{\zeta}{1-\Lambda_{\min}} \cdot D[o,\ell_o](t_o) \end{array} \tag{3}$$

Finally, we upper-bound the free-flow distance of ℓ_o from d by exploiting the triangle inequality:

$$\begin{array}{rcl} \underline{D}[\ell_o,d] & \leq & D[\ell_o,d](t_o - x_o) \leq D[\ell_o,o](t_o - x_o) + D[o,d](t_o) \\ & \overset{/*\ (3)\ */}{\leq} & \frac{\zeta}{1-\Lambda_{\min}} \cdot D[o,\ell_o](t_o) + D[o,d](t_o) \end{array}$$

which is exactly the desired inequality. ◀

The following lemma provides an upper bound on the approximation guarantee of HQA, when an (ESC)-termination occurs.

▶ **Lemma 16.** *Assume that $\ell_o \in L$ is informed (i.e., $d \in C[\ell_o]$) and settled by the initial* TDD *ball that* HQA *grows from (o, t_o). Then, for any given value $\varphi \geq 1$, if* HQA *terminated due to occurrence of (ESC), the reported travel-time is at most a $\left(1 + \varepsilon + \frac{\psi}{\varphi \cdot (r+1)}\right)$-approximation of the minimum travel-time $R_d = D[o,d](t_o)$.*

Proof of Lemma 16. Let $D[o,\ell_o](t_o) = R_o$. Then, $R_d \leq \overline{\Delta}[o,d](t_o) = R_o + \overline{\Delta}[\ell_o,d](t_o+R_o)$. By Theorem 14 (cf. also the analysis of `FCA` in [16]) we can easily deduce that:

$$\begin{aligned}
& \frac{\overline{\Delta}[\ell_o,d](t_o+R_o)}{R_o} \leq (1+\varepsilon)\frac{R_d}{R_o} + \psi - 1 \\
\overset{\text{/* } (ESC)\text{-termination */}}{\Rightarrow} \quad & (1+\varepsilon)\cdot\varphi\cdot(r+1) + \psi - 1 \leq \frac{\overline{\Delta}[\ell_o,d](t_o+R_o)}{R_o} \leq (1+\varepsilon)\frac{R_d}{R_o} + \psi - 1 \\
\Rightarrow \quad & \frac{R_d}{R_o} \geq \varphi\cdot(r+1) \\
\Rightarrow \quad & 1+\varepsilon+\frac{\psi R_o}{R_d} \leq 1+\varepsilon+\frac{\psi}{\varphi\cdot(r+1)}
\end{aligned}$$

Since $1+\varepsilon+\frac{\psi R_o}{R_d}$ is an upper bound on the approximation guarantee provided by `FCA` (cf. Theorem 14), which is indeed simulated by `HQA` until the determination of the appropriate level in the hierarchy, we conclude that the eventual solution that will be provided by `HQA` is at least as good, since the (ESC)-termination returns the best approximate solution seen so far via an informed landmark, among which is also the one that goes via ℓ_o. ◀

We proceed now by studying the first appearance of a level-i landmark within the unique outgoing ball from (o,t_o). The next lemma shows that, with high probability, this first appearance of a level-i landmark will take place in the following *ring* for level-i:

$$\begin{aligned}
RING[o;i](t_o) \quad &:= \quad B\left[o\,;\ (N_i)^{\delta/(r+1)}\cdot\ln(n)\right](t_o) \ \setminus\ B\left[o\,;\ \frac{(N_i)^{\delta/(r+1)}}{\ln(n)}\right](t_o) \\
&= \quad B\left[o\,;\ \frac{\ln(n)}{\rho_i}\right](t_o) \ \setminus\ B\left[o\,;\ \frac{1}{\rho_i\ln(n)}\right](t_o)
\end{aligned}$$

since $N_i = n^{(\gamma^i-1)/\gamma^i}$ and $\rho_i = n^{-\delta\cdot(\gamma^i-1)/[(r+1)\gamma^i]} = N_i^{-\delta/(r+1)}$.

▶ **Lemma 17.** *$\forall i \in [k]$, there is at least one level-i landmark in $RING[o;i](t_o)$, with probability $1-\mathcal{O}\left(\frac{1}{n}\right)$.*

Proof of Lemma 17. Consider any subset of vertices $S \subseteq V$, of size $s = |S| \in \mathbb{N}$. The probability that none of the vertices in S is a level-i landmark (i.e., from L_i) is $(1-\rho_i)^s \leq \exp(-s\rho_i)$.

Observe now that, for $i \in [k]$, $s_i = |RING[o;i](t_o)| = \frac{\ln(n)}{\rho_i} - \frac{1}{\rho_i\ln(n)}$. Thus, we conclude that: $\mathbb{P}\left\{|RING[o;i](t_o)\cap L_i| = 0\right\} \leq \exp(-s_i\cdot\rho_i) = \frac{\exp\left(\frac{1}{\ln(n)}\right)}{n} \in \mathcal{O}\left(\frac{1}{n}\right)$. ◀

The next lemma states that, given the actual Dijkstra-Rank $\Gamma[o,d](t_o)$ that we seek for, the level-i landmark (if any) that is settled within $RING[o;i](t_o)$ is indeed informed about d, for all i such that $N_i \geq \Gamma[o,d](t_o)$.

▶ **Lemma 18.** *For $i \in [k]$, let $\Gamma[o,d](t_o) \leq N_i = n^{\frac{\gamma^i-1}{\gamma^i}}$. Assuming that (ALH)-termination occurred, the first level-i landmark $\ell_{i,o} \in L_i \cap RING[o;i](t_o)$ that is settled by the initial `TDD` ball grown by `HQA`, has $d \in C[\ell_{i,o}]$.*

Proof of Lemma 18. Let $R_{i,o} = D[o,\ell_{i,o}](t_o)$ and recall that $R_d = D[o,d](t_o)$. Assume also that $R_d > R_{i,o}$, because otherwise an exact solution will be anyway discovered before $\ell_{i,o}$ is

settled. Then we have, by Lemma 15:

$$\begin{aligned}
\underline{D}[\ell_{i,o}, d] &\leq \frac{\zeta}{1-\Lambda_{\min}} R_{i,o} + R_d \\
&< \left(1 + \frac{\zeta}{1-\Lambda_{\min}}\right) \cdot R_d \\
&\overset{/* \text{ As. 2.4 } */}{\leq} \left(1 + \frac{\zeta}{1-\Lambda_{\min}}\right) \cdot g(n) \cdot n^{(\gamma^i - 1)/(\lambda\gamma^i)} \\
\Rightarrow \underline{\Gamma}[\ell_{i,o}, d] &\overset{/* \text{ As. 2.4 } */}{\leq} f(n) \cdot (\underline{D}[\ell_{i,o}, d])^\lambda \leq f(n) \cdot g^\lambda(n) \cdot \left(1 + \frac{\zeta}{1-\Lambda_{\min}}\right)^\lambda \cdot n^{(\gamma^i-1)/\gamma^i} \\
&= n^{(1+\lambda)\mathcal{O}\left(\frac{\log\log(n)}{\log(n)}\right)} \cdot n^{\lambda \log\left(1+\frac{\zeta}{1-\Lambda_{\min}}\right)/\log(n)} \cdot n^{(\gamma^i-1)/\gamma^i} \\
&= n^{(\gamma^i-1)/\gamma^i + o(1)} \leq n^{(\gamma^i-1)/\gamma^i + \xi_i}
\end{aligned}$$

for any $\xi_i \geq \frac{(1+\lambda)\cdot\log\log(n) + \lambda\log\left(1+\frac{\zeta}{1-\Lambda_{\min}}\right)}{\log(n)}$ which is certainly true since $\lambda \in o\left(\frac{\log(n)}{\log\log(n)}\right)$ and $f(n), g(n) \in \text{polylog}(n) = n^{\mathcal{O}\left(\frac{\log\log(n)}{\log(n)}\right)}$. The last inequality implies that $d \in C[\ell_{i,o}]$. ◀

The next lemma states that at no level-j in the hierarchy earlier than the *appropriate level* i^* corresponding to $\Gamma[o,d](t_o)$ (cf. Theorem 12), may `HQA` find a landmark which contains d in its coverage (and which would then cause an incorrect "guess" of the appropriate level for the query at hand), provided that no (ESC)-termination occurred.

▶ **Lemma 19.** *For $i \geq 1$, assume that $\Gamma[o,d](t_o) > N_i = n^{(\gamma^i-1)/\gamma^i}$ and, while growing the ball from (o, t_o), no (ESC) occurs. Then, $\forall 1 \leq j \leq i$ no level-j landmark ℓ in $RING[o;j](t_o)$ contains d in its coverage: $d \notin \cup_{1\leq j\leq i} \cup_{\ell \in L_j \cap RING[o;j](t_o)} C[\ell]$.*

Proof of Lemma 19. Recall again that $R_{j,o} = D[o, \ell_{j,o}](t_o)$ is the distance from the origin o to the first level-j landmark that we meet in $RING[o;j](t_o)$. We start by providing a simple proof for the cases of $1 \leq j \leq i-1$. We shall then handle the case $j = i$ separately, since it is a little bit more involved. So, fix arbitrary $j \in [i-1]$. For any $\ell_{j,o} \in RING[o;j](t_o) \cap L_j$ it holds, by the triangle inequality, that:

$$\begin{aligned}
&D[\ell_{j,o}, d](t_o + R_{j,o}) \\
&\quad\geq D[o,d](t_o) - D[o, \ell_{j,o}](t_o) \\
&\overset{/* \text{ As. 2.4 } */}{\geq} \frac{1}{f^{1/\lambda}(n)} (\Gamma[o,d](t_o))^{1/\lambda} - g(n)(\Gamma[o,\ell_{j,o}](t_o))^{1/\lambda} \\
&\quad> \frac{1}{f^{1/\lambda}(n)} n^{(\gamma^i-1)/(\lambda\gamma^i)} - g(n)\ln^{1/\lambda}(n) n^{a(\gamma^j-1)/((r+1)\lambda\gamma^j)} \\
&\quad= n^{(\gamma^j-1)/(\lambda\gamma^j)} \cdot \left[\frac{n^{(\gamma^{i-j}-1)/(\lambda\gamma^i)}}{f^{1/\lambda}(n)} - g(n)\frac{\ln^{1/\lambda}(n)}{n^{\left(1-\frac{a}{r+1}\right)(\gamma^j-1)/(\lambda\gamma^j)}}\right]
\end{aligned}$$

$$\begin{aligned}
\Rightarrow \Gamma[\ell_{j,o}, d](t_o + R_{j,o}) &\overset{/* \text{ As. 2.4 } */}{\geq} \left(\frac{D[\ell_{j,o}, d](t_o + R_{j,o})}{g(n)}\right)^\lambda \\
&> n^{(\gamma^j-1)/(\gamma^j)} \cdot \left[\frac{n^{(\gamma^{i-j}-1)/(\lambda\gamma^i)}}{g^\lambda(n) f^{1/\lambda}(n)} - \frac{\ln^{1/\lambda}(n)}{g^{\lambda-1}(n) n^{\left(1-\frac{\delta}{r+1}\right)(\gamma^j-1)/(\lambda\gamma^j)}}\right]^\lambda
\end{aligned}$$

Observe now that, from this last inequality and Assumption 2.3, the following is deduced:

$$\begin{aligned}\underline{\Gamma}[\ell_{j,o},d] &> \frac{n^{(\gamma^j-1)/(\gamma^j)}}{\text{polylog}(n)}\cdot\left[\frac{n^{(\gamma^{i-j}-1)/(\lambda\gamma^i)}}{g^\lambda(n)f^{1/\lambda}(n)}-\frac{\ln^{1/\lambda}(n)}{g^{\lambda-1}(n)n^{(1-\frac{\delta}{r+1})(\gamma^j-1)/(\lambda\gamma^j)}}\right]^\lambda\\ &\geq n^{\frac{\gamma^j-1}{\gamma^j}+\frac{\gamma^{i-j}-1}{\gamma^i}-\mathcal{O}\left(\lambda^2\cdot\frac{\log\log(n)}{\log(n)}\right)}\\ &> n^{(\gamma^j-1)/(\gamma^j)+\xi_j}\end{aligned}$$

The last inequality holds when $\lambda\in o\left(\sqrt{\frac{\log(n)}{\log\log(n)}}\right)$, and $\xi_j<\gamma^{-j}-\gamma^{-i}=\frac{\gamma-1}{\gamma^{j+1}}$. The last inequality implies also that $d\notin C[\ell_{j,o}]$.

We shall now study separately the case $j=i$. Apart from the containment of the level-i landmarks in $RING[o;i](t_o)$, we also exploit the fact that the first of these landmarks met by the unique TDD ball from (o,t_o) does not cause an (ESC)-termination of HQA. This implies that:

$$\begin{aligned}\frac{\overline{\Delta}[\ell_{i,o},d](t_o+R_{i,o})}{R_{i,o}} &< (1+\varepsilon)\beta(r+1)+\psi-1=:\chi\\ \Rightarrow R_{i,o} &> \frac{\overline{\Delta}[\ell_{i,o},d](t_o+R_{i,o})}{\chi}\geq\frac{D[\ell_{i,o},d](t_o+R_{i,o})}{\chi}\\ &\overset{/*\text{ triangle ineq. }*/}{\geq} \frac{R_d-R_{i,o}}{\chi}\\ \Rightarrow R_{i,o} &> \frac{R_d}{1+\chi}\\ &\overset{/*\text{ As. 2.4 }*/}{>} \frac{n^{(\gamma^i-1)/(\lambda\gamma^i)}}{(1+\chi)f^{1/\lambda}(n)}\end{aligned}\tag{4}$$

Nevertheless, we also know that for any $\ell_{i,o}\in RING[o;i](t_o)$ the following holds:

$$\begin{aligned}&\Gamma[o,\ell_{i,o}](t_o) \leq n^{\delta(\gamma^i-1)/((r+1)\gamma^i)}\cdot\ln(n)\\ \overset{/*\text{ As. 2.4 }*/}{\Rightarrow}\quad R_{i,o}=D[o,\ell_{i,o}](t_o) &\leq g(n)\left(\Gamma[o,\ell_{i,o}](t_o)\right)^{1/\lambda}\\ &\leq g(n)n^{\delta(\gamma^i-1)/((r+1)\lambda\gamma^i)}\ln^{1/\lambda}(n)\\ &\leq \frac{n^{(\gamma^i-1)/(\lambda\gamma^i)}}{(1+\chi)f^{1/\lambda}(n)}\end{aligned}\tag{5}$$

Inequality (5) holds if and only if

$$\begin{aligned}n^{\left(1-\frac{\delta}{r+1}\right)(\gamma^i-1)/(\lambda\gamma^i)} &\geq (1+\chi)g(n)f^{1/\lambda}(n)\ln^{1/\lambda}(n)\\ \Leftrightarrow\left(1-\frac{\delta}{r+1}\right)\frac{\gamma^i-1}{\gamma^i}\log(n)-\log(\ln(n))-\log(f(n)) &\geq \lambda\log(1+\chi)+\lambda\log(g(n))\\ \Leftrightarrow\left(1-\frac{\delta}{r+1}\right)\frac{\gamma^i-1}{\gamma^i}\log(n)-\mathcal{O}(\log\log(n)) &\geq \lambda\left[\log(1+\chi)+\mathcal{O}(\log\log(n))\right]\end{aligned}$$

which is certainly true for $\lambda\in o\left(\frac{\log(n)}{\log\log(n)}\right)$ and $f(n),g(n)\in\text{polylog}(n)$. From (4) and (5) we are led to a contradiction. Therefore, any level-i landmark in $RING[o;i](t_o)$ either does not possess a travel-time summary for d, or causes an early-stopping of HQA upon its settlement. ◀

The last lemma proves that, given that an (ALH)-termination occurs, HQA achieves the desired approximation guarantee.

▶ **Lemma 20.** *For $i^* \in \{2,3,\ldots,k+1\}$, let $N_{i^*-1} < \Gamma[o,d](t_o) \leq N_{i^*}$. Assume also that an (ALH)-termination occurred. Then, an $\left(1+\varepsilon\frac{(1+\varepsilon/\psi)^{r+1}}{(1+\varepsilon/\psi)^{r+1}-1}\right)$−approximate solution is returned by* `HQA`*, with probability $1-\mathcal{O}\left(\frac{1}{n}\right)$. The expectation of the query-time is $\mathcal{O}((N_{i^*})^a \operatorname{polylog}(n))$.*

Proof of Lemma 20. As it was already explained by Lemmata 17, 18 and 19, we know that a successful "guess" of i^* will occur with probability $1-\mathcal{O}\left(\frac{1}{n}\right)$. Then, `HQA` proceeds with the execution of $\texttt{RQA}_{i^*}$, whose expected time is $\mathcal{O}((N_{i^*})^a \operatorname{polylog}(n))$ (cf. Theorem 14).

As for the approximation guarantee, since the analysis for $\texttt{RQA}_{i^*}$ is based solely on the quality of the paths via landmarks discovered from ball centers located along a minimum-travel-time *od*-path $p^* \in SP[o,d](t_o)$ (cf. corresponding proof for `RQA` in [16]), it suffices to prove that all the level-$(\geq i^*)$ landmarks discovered from ball centers which reside at the (unknown) shortest *od*-path, are informed about d.

For an arbitrary ball center $w_{i^*,j} \in p^*$ and its closest level-i landmark for some $i \geq i^*$, $\ell_{i,j} \in L_i$, let $R_{i^*,j} = D[w_{i^*,j},\ell_{i,j}](t_o + D[o,w_{i^*,j}](t_o))$. Then, either $R_d \leq R_{i^*,o} + \ldots + R_{i^*,j}$, in which case an exact solution is returned, or else the following holds:

$$
\begin{aligned}
\underline{D}[\ell_{i,j},d] \quad &\leq \quad \frac{\zeta}{1-\Lambda_{\min}} R_{i^*,j} + R_d < \left(1+\frac{\zeta}{1-\Lambda_{\min}}\right)\cdot R_d \\
&\overset{/*\text{ As. 2.4 }*/}{\leq} \quad \left(1+\frac{\zeta}{1-\Lambda_{\min}}\right)\cdot g(n)\cdot n^{(\gamma^i-1)/(\lambda\gamma^i)} \\
\Rightarrow \underline{\Gamma}[\ell_{i,j},d] \quad &\overset{/*\text{ As. 2.4 }*/}{\leq} \quad f(n)\cdot(\underline{D}[\ell_{i,j},d])^\lambda \leq f(n)\cdot\left(1+\frac{\zeta}{1-\Lambda_{\min}}\right)^\lambda \cdot g^\lambda(n)\cdot n^{(\gamma^i-1)/\gamma^i} \\
&\leq \quad n^{(\gamma^i-1)/\gamma^i+\xi_i}
\end{aligned}
$$

for $\lambda\log(g(n)) + \log(f(n)) \leq \xi_i \log(n) \Leftrightarrow \xi_i \geq (\lambda+1)\cdot\mathcal{O}\left(\frac{\log\log(n)}{\log(n)}\right)$. Again, a sufficient condition for this is that $\lambda \in \mathrm{o}\left(\frac{\log(n)}{\log\log(n)}\right)$. We therefore conclude that $d \in C[\ell_{i,j}]$.

Since all the discovered landmarks from ball centers along the unknown shortest path are informed, the analysis of $\texttt{RQA}_{i^*}$ provides the claimed approximation guarantee. ◀

Having proved all the required building blocks, we are now ready to provide the proof of our main argument for `HORN`.

Proof of Theorem 12. We start with a short sketch of the proof, before delving into the details: The claimed preprocessing space and time are proved, mainly due to the appropriate choice of the ξ_i parameters. We thus focus on the more involved analysis of the time-complexity of `HQA`. First of all, exploiting Assumptions 2.1 and 2.2, we prove an upper bound on the free-flow travel-time of the destination d from an arbitrary landmark ℓ, as a function of the (unknown) travel-time $D[o,d](t_o)$. We then prove that if the (ESC)-termination criterion occurs, then we already have a satisfactory approximation guarantee. Otherwise, assuming that the Dijkstra Rank $\Gamma[o,d](t_o)$ lies within the two consecutive values N_{i^*-1} and N_{i^*}, we exploit the triangle inequality along with Assumptions 2.3 and 2.4, to prove that, with high probability, the closest level-i^* landmark $\ell_{i^*,o}$ to the origin o lies in a certain ring (i.e., difference of two cocentric balls from o) and is indeed "*informed*". Additionally, all the landmarks discovered prematurely with respect to the ring of the corresponding level to which they belong, cannot be "*informed*", given the failure of the (ESC)-termination criterion. Moreover, discovered landmarks of higher levers, despite being informed, do not lie in the ring of the corresponding levels and therefore do not interrupt the guessing procedure. This then implies that the appropriate variant of `RQA` will be successfully executed and will

return a travel-time estimation with the required approximation guarantee and query-time complexity sublinear in N_{i^*} (the targeted Dijkstra-Rank).

We now proceed with the detailed presentation of the proof. For each level $i \in [k]$, the targeted Dijkstra-Rank is $N_i = n^{(\gamma-1)/\gamma^i}$. The coverage of each level-i landmark contains $c_i = N_i \cdot n^{\xi_i} = n^{(\gamma-1)/\gamma^i+\xi_i}$ destinations, $F_i = (c_i)^{\theta_i/\nu}$ of which are discovered by `BIS` and the remaining $c_i - F_i$ destinations are discovered by `TRAP`.

We start with the analysis of the approximation guarantee and the query-time complexity for `HQA`. Observe that the query-time is dominated by the scenario in which there is an (ALH)-termination (the discovery of the destination, or an (ESC)-termination will only improve the performance of the algorithm). Assume that for some $i^* \in \{2, 3, \ldots, k+1\}$ (which we call the *appropriate level* for the query at hand), $N_{i^*-1} = n^{\frac{\gamma^{i^*-1}-1}{\gamma^{i^*-1}}} < \Gamma[o,d](t_o) \leq n^{\frac{\gamma^{i^*}-1}{\gamma^{i^*}}} = N_{i^*}$. Lemma 16 assures that, when an (ESC)-termination occurs, an $\left(1+\varepsilon+\frac{\psi}{\varphi\cdot(r+1)}\right)$-approximate solution was anyway discovered and thus there is no need to guarantee the success of the guessing procedure. By setting $\varphi = \frac{\varepsilon(r+1)}{\psi(1+\varepsilon/\psi)^{r+1}-1}$, the approximation guarantee when (ESC)-termination occurs is equal to $1+\varepsilon\frac{(1+\varepsilon/\psi)^{r+1}}{(1+\varepsilon/\psi)^{r+1}-1}$. Lemma 17 proves that, with probability $1-\mathcal{O}\left(\frac{1}{n}\right)$, at least one appearance of a level-i^* landmark occurs in the appropriate ring: $L_{i^*} \cap RING[o;i^*](t_o) \neq \emptyset$. Lemma 18 then proves that, given that an (ALH)-termination occurs, the first level-i^* landmark $\ell_{i^*,o}$ settled by the initial `TDD` ball within $RING[o;i^*](t_o)$ is indeed an "informed" landmark: $d \in C[\ell_{i^*,o}]$. Lemma 19 proves that, given an (ALH)-termination, no landmark of a previous level $j < i^*$ that was settled before $\ell_{i^*,o}$ may be informed: $\forall \ell \in \cup_{j\in[i^*-1]}(L_j \cap RING[o;j](t_o)),\ d \notin C[\ell]$. As for landmarks $\ell \in \cup_{j=i^*+1}^{k+1} L_j$, if we settle ℓ before $\ell_{i^*,o}$ and it happens that $d \in C[\ell]$ (i.e., ℓ is "informed"), this event will not interrupt the guessing procedure of `HQA`, because it is not in the ring of the corresponding level: $(B[o](t_o) \cap \cup_{j=1}^{i^*} RING[o;j](t_o)) \cap (\cup_{j=i^*+1}^{k+1} RING[o;j](t_o)) = \emptyset$. Therefore we conclude that, in the case of an (ALH)-termination, with probability $1-\mathcal{O}\left(\frac{1}{n}\right)$, the "guess" of the appropriate level-i^* is indeed correct. Finally, Lemma 20 demonstrates that, when an (ALH)-termination occurs with a successful "guess" of i^*, a $\left(1+\varepsilon\frac{(1+\varepsilon/\psi)^{r+1}}{(1+\varepsilon/\psi)^{r+1}-1}\right)$-approximate solution is returned by `HQA`.

As for the expected query-time of `HQA`, we wish to assure that it is sublinear in the Dijkstra-Rank of the appropriate level i^* of the hierarchy. Suppose that the exponent of sublinearity is $\delta \in (a,1)$, i.e., the query-time is comparable to $\mathcal{O}\left((N_{i^*})^\delta\right)$, where $a \in (0,1)$ is the exponent relating the period of the metric with the network size ($T = n^a$). We focus on the highly-probable event that `HQA`, if it terminates due to the (ALH) criterion, makes a correct guess of i^*. Conditioned on this event, the expected cost of $\texttt{RQA}_{i^*}$ gives the appropriate value for the exponent ω_i of the landmark-sampling probability:

$$\left(\frac{1}{\rho_{i^*}}\right)^{r+1} \log\left(\frac{1}{\rho_{i^*}}\right) \text{polylog}(n) = n^{\omega_i(r+1)+o(1)} = (N_{i^*})^\delta \Rightarrow \boxed{\omega_i = \frac{\delta}{r+1}\frac{\gamma^{i^*}-1}{\gamma^{i^*}}}$$

The contribution to the expected query-time of `HQA` of the unlikely event that the algorithm makes a wrong guess about the appropriate level of the query at hand is negligible, due to the quite small probability of this happening.

We proceed next with the study of the required time and space for the `BIS` + `TRAP`-based preprocessing of the `HORN` oracle. We wish to bound the preprocessing requirements with $(k+1)\cdot n^{2-\beta}$, for a given $\beta > 0$. We shall make the appropriate choices of our tuning parameters so that in each level it holds that the preprocessing requirements are $S_i\ ,\ P_i \leq n^{2-\beta}$.

We begin with the determination of the requirements of level-i, for each $i \in [k]$. For this level we generate $|L_i| = n^{1-\omega_i}$ landmarks, each of which possesses travel-time summaries for all the c_i destinations contained in its own coverage. Based on the analysis of the preprocessing analysis of FLAT (cf. Theorem 9), but now restricting ourselves within the coverage of each landmark, we know that the overall preprocessing requirements of level-i are bounded as follows:

$$\begin{aligned} S_i \ , \ P_i &\in \mathcal{O}\big(|L_i| \cdot \big(F_i^2 \operatorname{polylog}(F_i) + T^{1-\theta_i} c_i\big)\big) \\ &= \mathcal{O}\left(n^{1-\omega_i} \cdot \left[n^{\frac{2\theta_i}{\nu}\left(\frac{\gamma^i-1}{\gamma^i}+\xi_i\right)+\alpha(1)} + n^{a(1-\theta_i)+\frac{\gamma^i-1}{\gamma^i}+\xi_i}\right]\right) \\ &\subseteq \mathcal{O}\left(n^{1-\omega_i} \cdot \left[n^{\frac{2\theta_i}{\nu}+\alpha(1)} + n^{a(1-\theta_i)+1}\right]\right) \end{aligned}$$

provided that $\xi_i \leq \gamma^{-i}$. We choose again $\theta_i = \frac{1+a}{2/\nu+a}$, which in turn assures that

$$S_i \ , \ P_i \in n^{2-\omega_i+a(1-\theta_i)+\alpha(1)}$$

We correlate our demand for subquadratic preprocessing with the recursion budget (and thus, the approximation guarantee) that we can achieve:

$$\begin{aligned} & 2-\omega_i+a\cdot(1-\theta_i) < 2-\beta < 2 \\ \Leftrightarrow \quad & \beta < \omega_i - a\cdot(1-\theta_i) = \tfrac{\delta(1-\gamma^{-i})}{r+1} - a\cdot\tfrac{2/\nu-1}{2/\nu+a} \\ \Leftrightarrow \quad & \boxed{r < \tfrac{\delta}{a}\cdot\tfrac{(2/\nu+a)(1-\gamma^{-i})}{\beta\cdot(2/(a\nu)+1)+2/\nu-1} - 1} \end{aligned}$$

Therefore, the overall preprocessing requirements of HORN are $S_{\texttt{HORN}} \ , \ P_{\texttt{HORN}} \in (k+1)\cdot n^{2-\beta} = n^{2-\beta+\frac{\log(k+1)}{\log(n)}} = n^{2-\beta+\alpha(1)}$ for $k \in \operatorname{polylog}(n)$. ◀